

Detecting causal associations in large nonlinear time series datasets

Jakob Runge^{1,*}, Dino Sejdinovic^{2,3}, and Seth Flaxman²

1 Grantham Institute, Imperial College, London SW7 2AZ, United Kingdom

2 Department of Statistics, University of Oxford, Oxford OX1 3LB, United Kingdom

3 The Alan Turing Institute for Data Science, London NW1 3DB, United Kingdom

* jakobrunge@posteo.de

Abstract

Detecting causal associations in time series datasets is a key challenge for novel insights into complex dynamical systems such as the Earth system or the human brain. Interactions in high-dimensional dynamical systems often involve time-delays, nonlinearity, and strong autocorrelations. These present major challenges for causal discovery techniques such as Granger causality leading to low detection power, biases, and unreliable hypothesis tests. Here we introduce a reliable and fast method that outperforms current approaches in detection power and scales up to high-dimensional datasets. It overcomes detection biases, especially when strong autocorrelations are present, and allows ranking associations in large-scale analyses by their causal strength. We provide mathematical proofs, evaluate our method in extensive numerical experiments, and illustrate its capabilities in a large-scale analysis of the global surface-pressure system where we unravel spurious associations and find several potentially causal links that are difficult to detect with standard methods. The broadly applicable method promises to discover novel causal insights also in many other fields of science.

Introduction

How do coupled subprocesses in the Earth’s atmosphere-ocean system bring about a large-scale climate phenomenon like El Niño Southern Oscillation or the North Atlantic Oscillation (NAO) with their profound impact on global climate? Through which pathways are different brain regions interacting to make the mind work? How does the Sun influence the Earth’s magnetosphere causing satellites and power grid infrastructure to fail? Progress towards an improved analytical understanding of such complex dynamical systems is an interdisciplinary challenge requiring better modeling as well as data science advancements. Regarding the latter, the roots of our limitations are partly due to data constraints, but also due to a lack of reliable methods for extracting *causal* associations from large-scale datasets that go beyond purely *pairwise* association analyses [6, 41, 42].

Towards a network more faithfully representing causal associations (Fig. 1), advanced methodologies such as Granger causality (GC) [16], dynamical systems state-space causality [54], dynamical Bayesian networks [27], or the framework of causal discovery algorithms such as the PC algorithm [34, 52] (named after its inventors Peter Spirtes and Clark Glymour) have been proposed. While such approaches do not provide causal information at the level of laboratory or field experiments, these are not feasible in such large complex systems as the Earth. Then causal discovery can help in formulating and testing hypotheses to build and improve theories, simulation models, and, where possible, conceive new natural experiments.

Complex dynamical systems present major challenges for causal discovery methods due to issues such as high dimensionality, nonlinearity, time-delayed interactions, and strong autocorrelations. High dimensionality arises even for datasets of just a few variables because often large time lags have to be accommodated. This leads to low detection power for GC as we show here, and makes state-space [54] or score-based Bayesian network approaches computationally difficult to treat [27]. While the dimensionality of the problem can be




Figure 1. Problem setting and proposed causal discovery method. (A) Consider a high-dimensional, autocorrelated time series dataset with linear and nonlinear, time-delayed causal dependencies among the variables. (B) The goal of causal discovery is to reconstruct the causal network of these dependencies from the time series alone (labels denote time lags, node colors lag-1 autocorrelation strength, link colors causal link strength) and unveil spurious associations (gray arrows) which necessarily emerge due to common drivers or transitive indirect paths. Correlation matrices are, therefore, often very dense, while causal networks are typically sparse. (C) Since dependencies are typically time-delayed, the actual goal of causal discovery is to estimate the *time series graph* [13, 46], where lagged variables up to τ_{\max} are considered (see definitions in Sect. S1). Granger causality is estimated from the set of all lagged variables with an often very high dimension $N\tau_{\max}$, while our causal discovery technique avoids high dimensionality by a two-step method: (1) the much smaller set of parents \mathcal{P} of a variable (e.g., here $\mathcal{P}(X_t^1) = \{X_{t-1}^1, X_{t-1}^2\}$, blue boxes) is approximately estimated, and, (2) these parents are used as conditions in the MCI test. For the example marked by black dots, $\mathcal{P}(X_t^1)$ (blue boxes) denotes the conditions necessary to establish conditional independence, while the additional condition on the parents $\mathcal{P}(X_{t-2}^3)$ (red boxes) eliminates autocorrelation and other biases in the MCI test.

large, causal interdependencies are typically very sparse in comparison and are necessarily vastly outnumbered by spurious associations emerging from indirect interactions or common drivers (Fig. 1B). Causal discovery algorithms such as the PC algorithm [46, 51] exploit this property by an iterative approach but suffer from unreliable hypothesis tests as we show here. In addition, we will see that strong autocorrelations in time series lead to inflated false detections and undesirable biases. The detection probability can depend, for example, not only on causal strength but also on how strongly autocorrelated the variables are. In summary, these problems lead to brittle causal discoveries and a more reliable methodology suited for time series is required before a wider adoption of causal discovery algorithms in complex dynamical systems.

We present a causal discovery method suitable for large-scale time series datasets featuring linear as well as nonlinear, time-delayed dependencies. It is based on a conditional independence test for time series that, firstly, outperforms alternative autocorrelation remedies (such as block-shuffling or *pre-whitening*) yielding reliable rates of false positives that match the prescribed significance level even in the presence of strong autocorrelations. Secondly, our test largely eliminates detection biases from strong autocorrelations and other drivers leading to an improved ‘causal signal-to-noise ratio’. Finally, the method is scalable regarding computational complexity with higher detection power and comparably shorter runtimes than GC or the PC algorithm on large-scale time series datasets.

We provide mathematical proofs and compare our method in extensive numerical experiments to GC and the PC algorithm and illustrate its capabilities in a large-scale analysis of the global surface-pressure system. We unravel spurious associations and find several potentially causal links that are difficult to detect with standard methods, focusing in particular on drivers of the NAO.

Causal discovery method

Our method to estimate the causal *time series graph* [13, 46] (Fig. 1) is based on a two-step procedure we call PCMCI (see Sect. S1): (i) For every variable, an *approximate* estimate of the parents \mathcal{P} (defined in Fig. 1C) is

reconstructed using a fast variant of the PC algorithm (Algorithm S1). (ii) This low-dimensional set of parents is then used as a condition to perform a novel independence test adapted to autocorrelated time series (Fig. 1C, Algorithm S2) called *momentary conditional independence* (MCI). For a variable pair $(X_{t-\tau}^i, X_t^j)$ with time-delay τ , MCI is defined as

$$\text{MCI: } X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{P}(X_t^j) \setminus \{X_{t-\tau}^i\}, \mathcal{P}(X_{t-\tau}^i), \quad (1)$$

where ‘ $\perp\!\!\!\perp$ ’ denotes (conditional) independence. The MCI test can be implemented with different linear or nonlinear independence measures as discussed below. The condition on the parents $\mathcal{P}(X_t^j)$ (blue boxes in Fig. 1C) as used in the original PC algorithm [46] is sufficient to establish conditional independence between $X_{t-\tau}^i$ and X_t^j . But, as we show here, the additional conditioning on the parents of the supposed driver $X_{t-\tau}^i$ (red boxes) better controls false positives for strongly autocorrelated time series and makes MCI a valid criterion for causal strength. Then tests based on this criterion will have detection power that scales with causal strength only and is not biased by other factors.

As further elaborated on in Sect. S2, our notion of *causal strength* between variables $X_{t-\tau}^i$ and X_t^j at a lag $\tau > 0$ can be information-theoretically defined as the *momentary* information entering in X_t^j at time t (and not contained in the past of $X_{t-\tau}^i$), that is transferred to X_t^j [39, 43, 45]. Consider the model visualized in Fig. 1C with dependencies $X_t^1 = aX_{t-1}^1 + cX_{t-1}^2 + \eta_t^1$ and $X_t^2 = bX_{t-1}^2 + g(X_{t-1}^3) + \eta_t^2$, where η denotes independent noise and g an arbitrary smooth function. Here the causal strength between X_{t-1}^2 and X_t^1 only depends on the coefficient c , but, as we will see, the detection rate of the link in the independence test used in the original PC algorithm is strongly biased depending on the autocorrelation coefficients a, b and other common driver forcings.

For this example, the MCI test statistic I in its information-theoretic version can be derived (see Sect. S2) as $I(X_{t-1}^2; X_t^1 \mid \mathcal{P}(X_t^1) \setminus \{X_{t-1}^2\}, \mathcal{P}(X_{t-1}^2)) = I(\eta_{t-1}^2; c\eta_{t-1}^2 + \eta_t^1)$. This result also holds under more general dependencies as proven in Sect. S2 and implies that, firstly, if the noise terms η are independent, the MCI test is essentially performed on independent samples even in the presence of autocorrelations. Then common analytical or pre-computed null distributions apply and lead to a well-calibrated independence test. Secondly, MCI is an estimator of our notion of causal strength. Thus, the resulting MCI test has an improved causal signal-to-noise ratio allowing to rank causal links in large-scale studies in a meaningful way. For linear links, where causal strength can be attributed to a single coefficient c as in our example above, we prove (Sect. S2) that the MCI detection power is then unbiased regarding autocorrelations or common driver forcings (that can also be nonlinear). In Sect. S3 we show that both GC and the PC algorithm, on the other hand, are prone to biases due to autocorrelation and common drivers leading to a low causal signal-to-noise ratio and low detection power. Finally, in Sect. S2 we also prove that the two-step PCMCI method yields a consistent estimator of the causal time series graph.

Our method can be flexibly combined with any conditional independence test appropriate for the system under study. Here we implement two conditional independence test statistics (see Sect. S1): linear partial correlation (ParCorr) and a nonlinear test (GPACE) based on a Gaussian process (GP) regression [40] where the independence of the regression residuals is tested with a nonparametric alternating conditional expectation (ACE) estimator of the *maximal correlation* [5] which is computationally efficient and was found to have competitive statistical power [42]. Other choices include kernel conditional independence tests [15, 50, 60] or permutation tests based on mutual information [28].

Numerical experiments

We compare the performance of the proposed causal discovery method with GC and the PC algorithm as a standalone method (for implementation details see Sect. S1) on a general class of nonlinear stochastic models for a large variety of random dependence network topologies. The full model setup is detailed in Sect. S4, Fig. 2A gives an example network.

First, we investigate detection biases (Fig. 2). GC features expected false positive rates with a slight autocorrelation bias (Fig. 2B,C), but has low power, especially in the ParCorr implementation (Fig. 2B). Additionally, the power significantly depends on autocorrelation (see also Sect. S3 for analytical examples). The original PC algorithm is over-conservative in controlling false positives (except for some high autocorrelation

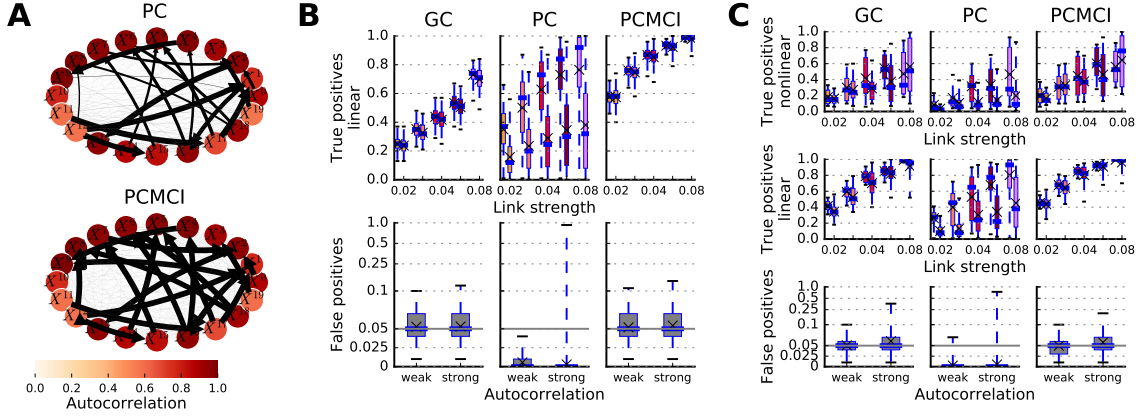


Figure 2. Numerical experiments demonstrating autocorrelation bias and scaling of power with link strength. The model setup is described in Sect. S4. **(A)** Example network with node color denoting autocorrelation strength, black arrow width proportional to true positive rate (TPR) and gray arrows to false positives (FPR, average across all lags). All true links have the same link strength. The top panel shows results for the original PC algorithm (Algorithm S1) that has very heterogeneous TPRs and FPRs. The bottom network shows results for the proposed PCMCI method with all true links having a higher and equal detection rate corresponding to their identical link strength. **(B)** Power-scaling property for linear ParCorr implementation (5% significance level). The pairs of boxplots (1%, 25%, 75%, 99% percentiles, median marked by a bar and mean marked by ‘x’) always denote weakly (left) and strongly (right) autocorrelated variable pairs. The bottom row shows the distribution of FPRs and the upper row the distributions of TPRs for different link strengths increasing along the x -axis in each plot. Note the logarithmic y -axis in the bottom panel for $FPR > 0.1$. The more confined a distribution, the less biased is the TPR. PCMCI shows no autocorrelation bias, controls FPRs at the expected rate, and has the highest power levels scaling well with link strength. **(C)** Same as **(B)**, but for models including linear as well as nonlinear dependencies analyzed with the GPACE implementation. Also here PCMCI more reliably controls FPRs and achieves the highest power levels. As proven in Sect. S2, an unbiased power (confined TPR distribution) can only be achieved for linear links where causal strength can be attributed to a single coefficient. See Figs. S6,S7,S8 for more results.

links) and has very heterogeneous power levels, strongly biased with links of the same causal strength but between higher autocorrelated variables having lower detection rates. The performance of the PC independence test for strong autocorrelations in high dimensions, also including alternative autocorrelation remedies, is investigated in more detail in another numerical experiment (Sect. S5 and Fig. S15).

PCMCI shows largely no autocorrelation bias, controls false positives at the expected rate, and has the highest power levels scaling well with link strength in both the ParCorr and GPACE implementation (strongly confined distributions in Figs. 2B,C). For nonlinear links (top row in Fig. 2C), the power levels in PCMCI are more heterogeneous as expected: A fully unbiased detection, independent of autocorrelation, is fundamentally possible only for linear causal links (while still allowing for nonlinear common drivers as analyzed in Sect. S2). Note that we use a small sample size here ($T = 250$), for larger sample sizes nonlinear links are much better detected (Figs. S8,S11,S12).

For larger network sizes with higher dimensionality, GC in the ParCorr implementation (Fig. 3A) has much lower detection power which is not much alleviated by regularized estimators. In the GPACE implementation (Fig. 3B) the power does not decrease as fast with the network size, but still the detection power is uniformly lower for GC compared to PCMCI and false positives are higher especially for high-autocorrelation links. PC has relatively stable but lower and more heterogeneous power levels than PCMCI, which has robustly high power even for network sizes with dimensions exceeding the sample size (e.g., $N_{\tau_{\max}} = 60 \cdot 5 = 300$ for $T = 150$ in the ParCorr implementation). Another advantage of PCMCI is the much shorter runtime for larger networks (top labels in Fig. 3). In the Supplement, we show that these results are robust for different parameter choices of the method and outperform pre-whitening as an autocorrelation remedy (Figs. S4,S5,S6,S7,S9,S10). The

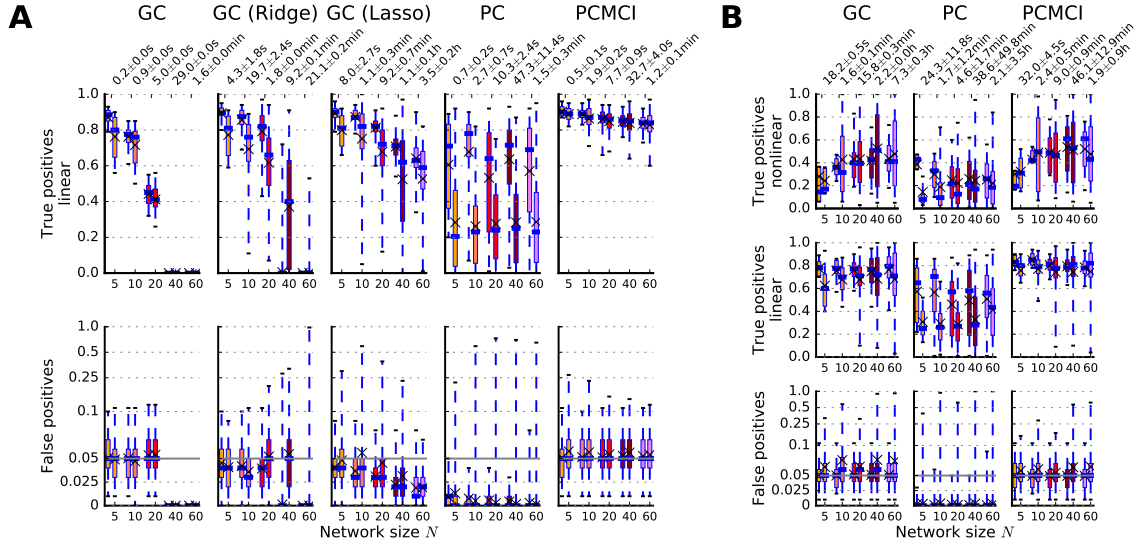


Figure 3. Testing high dimensionality with larger network sizes. All networks have $L = N$ randomly drawn links corresponding to an average degree of 1 (Sect. S4). **(A)** Linear models with ParCorr implementation. The pairs of boxplots (1%, 25%, 75%, 99% percentiles, median marked by a bar and mean marked by ‘x’) always denote weakly (left) and strongly (right) autocorrelated variable pairs. The bottom row shows the distribution of FPRs and the upper row the distributions of TPRs for different network sizes N increasing along the x -axis in each plot. Average runtime and standard deviation are given on top. PC shows over-conservative false positive control and low, heterogeneous power levels also for larger networks. While GC and PCMCI both well control FPRs, GC has steadily decreasing power levels and cannot be evaluated anymore for $N > 30$ since then the dimensionality exceeds the sample size. Regularized estimators using ridge and lasso regression (see Sect. S1) yield higher power levels, but also do not perform well in high dimensions and have much longer runtimes, while PCMCI avoids such high dimensionality and has rather robust power levels. **(B)** Same for GPACE implementation where GC exhibits a remarkable resistance to the curse of dimensionality, but still features higher FPRs, especially for high autocorrelation links, and lower power compared to PCMCI. Additionally, PCMCI has much shorter runtimes for larger networks. In Figs. S9,S10,S13,S14 we show more detailed results.

results also hold for larger sample sizes (Figs. S8,S11,S12), as well as denser causal network topologies with dependent noise terms (Figs. S13,S14).

Real data application

We demonstrate the method on a large-scale time-series dataset representing aspects of atmospheric dynamics in the Earth system where recent applications of causal discovery methods have already given promising results [12, 29, 48]. The dataset is described in [58] and consists of 60 surface-pressure component time series representing climatological subprocesses in the winters (December to February) of 1981–2012 at a 3-day time resolution. Taking into account lags of up to 3 weeks ($\tau_{\max} = 7$), the dimensionality of this dataset is $N_{\tau_{\max}} = 420$ with a comparably small sample size of about 950 samples and partially strong autocorrelations.

In Figs. 4A and B, we demonstrate how an exploratory causal discovery analysis based on the ParCorr implementation can unveil spurious associations and discover potential new causal mechanisms. Plotting the MCI test statistic values for each component pair against their Pearson correlations in Fig. 4A, we observe a large number of spurious links even though there are high correlation values. As further analyzed in Fig. S19, these components are spuriously correlated due to strong autocorrelations, common drivers, or due to the fact that they represent indirect links via other surface pressure components. The latter can be analyzed with causal mediation methods [47].

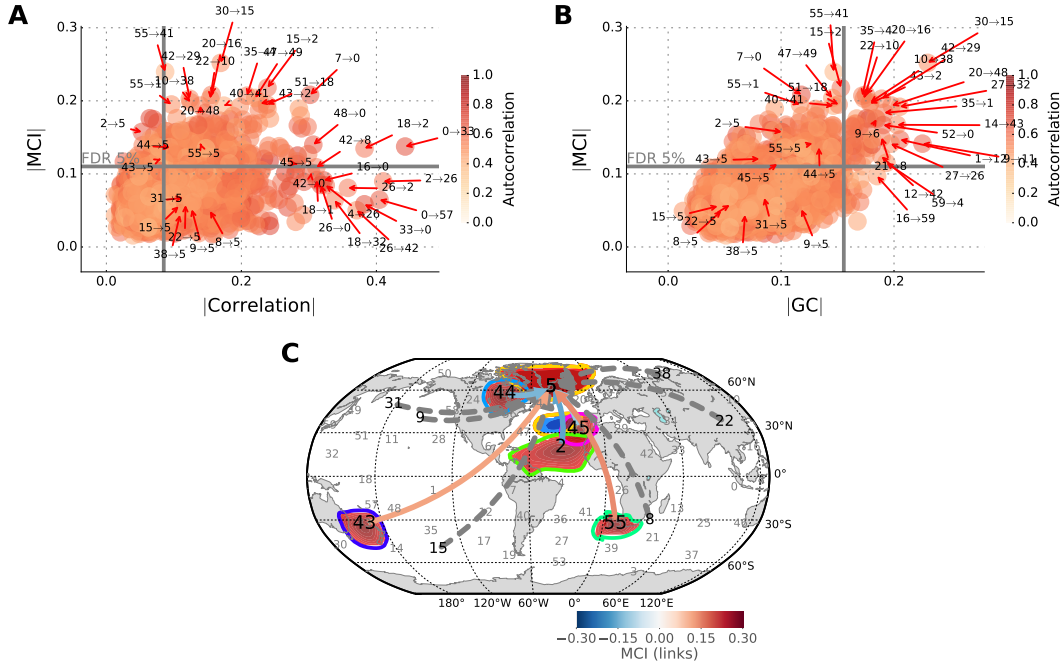


Figure 4. Real data example of the Earth’s surface-pressure system (see Sect. S7): (A) Comparison of MCI with Pearson correlation for each component pair: Scatter plot of value of $I_{i \rightarrow j}^{MCI}(\tau)$ at lag with maximum absolute MCI plotted against the Pearson correlation value (also at the corresponding lag with maximum absolute value). The point color depicts the mean of the two autocorrelations of a variable pair. With the gray horizontal and vertical lines defining significant links in each statistic, the lower right quadrant shows a large number of spurious correlations. On the other hand, links in the left upper half are causal links that would be overlooked by an exploratory correlation analysis. (B) Comparison of MCI with GC. Here the upper left quadrant shows a large number of causal links overlooked in a GC analysis due to its low power. In Tabs. S6,S7 we list p-values, Fig. S22 depicts the component regions for the strongest causal links and Figs. S19,S20 show comparisons with further measures. These results are robust against the method’s parameters (Fig. S21). (C) Drivers of component No. 5 representing the NAO: colored lines identify the different region loadings belonging to each causal driver component (note the dipole of the NAO, for all regions see [47]). Spurious correlation links are marked by gray dashed links. Further details and p-values are given in Fig. S23 and Tab. S5.

While Correlation is prone to detect spurious links, in Fig. 4B we show that the low power of GC leads to a large number of even very strong causal associations being overlooked. In contrast, the higher power and better interpretability of MCI as a measure of causal strength allows the ranking of links and to focus on the strongest causal relationships in large-scale causal discovery studies. In Fig. S22, we show examples of strong links in the pressure field and Tab. S6 lists p-values and results for GC and other measures. Some of the strong MCI links can be explained by well-known climatic mechanisms. For example, the anticorrelated links $30 \rightarrow 15$, $55 \rightarrow 41$ and $22 \rightarrow 10$ are likely to be surface pressure expressions of the downstream development of tropospheric Rossby wave patterns [19]. More long-range links such as $35 \rightarrow 4$ and $15 \rightarrow 2$ in the southern hemisphere and $30 \rightarrow 16$ in the northern hemisphere can at least partially be explained by Rossby waves sometimes propagating in upper-tropospheric wave guides [19]. Also, the time delay of a few days is in line with the typical speed of these wave guide mechanisms. For some links obtained in our analysis (e.g., $7 \rightarrow 0$ in Fig. S22), there is no obvious plausible explanation. Such exploratory findings can yield new hypotheses to be followed up by further data analyses involving more climatic variables (such as sea-surface temperatures) and also simulation experiments with weather and climate models in order to be rejected or confirmed and refined.

In Fig. 4C, we focus on drivers of the NAO represented here by the dipole component No. 5 (see also Fig. S23). The NAO is one of the most prominent patterns of atmospheric variability and strongly influences weather from eastern North America to Siberia [22]. A large number of components from the Pacific and the

Asian continent display significant correlations with the NAO. For example, the correlation of component No. 9 capturing the Pacific–North America pattern (PNA, but here with opposite sign) with the NAO is well-known in the climate literature [11] and it is an open question which mechanisms are behind it. Here we find that this correlation is spurious and can be explained by other surface-pressure drivers. A possible mechanism is an indirect link via one of the direct causal drivers of the NAO, for example, via a downstream Rossby wave development [59] through component No. 44 over East Canada as suggested in [11].

A pairwise correlation analysis is *mostly* prone to detect spurious links, but may also *overlook* causal links. Consider the strongly anticorrelated MCI link from the tropical Atlantic component No. 2 to No. 5. This link has only a very low Pearson correlation (Fig. 4A) and is also not detected in a GC analysis (Fig. 4B, see Tab. S5 for p-values). A climatological explanation may be anomalies in the local Hadley circulation (the overturning circulation in a latitude-height plane), associated with almost northward long Rossby wave propagation [20]. MCI better isolates this causal mechanism and reduces distorting effects from confounding drivers. In Sect. S3, we illustrate such a distorting influence in a model example: A common driver can, if of opposite sign, cancel out a strong correlation or GC value.

Discussion and Conclusion

Our flexible method allows accommodating a large variety of independence tests adapted to different types of data (see Sect. S1). For example, variables can be discrete or continuous, but networks can also be reconstructed with multivariate variables as single nodes in the graph. This flexibility can help to represent causal associations on different aggregation levels and opens up a way to study causal networks on multiple interdependent levels in the growing field of multilayer networks [4].

Our real data example illustrates that our causal discovery method is useful for theory building. Complemented by a subsequent causal pathway analysis [43, 47], it can help to analyze also indirect causal interaction mechanisms. Furthermore, causal drivers can also be utilized for optimal statistical predictions [44]. While we focus on time-lagged dependencies where there is no ambiguity about the causal direction, we note that a recently growing body of literature addresses the inference of causality if no time-ordering is given [32, 37, 52] which could be applied to determine causal directionality for contemporaneous links.

A fundamental assumption of any causal discovery method from observational data is *causal sufficiency* [52] implying that the term ‘causal’ has to be understood relative to the set of included variables. Non-included variables can still be the cause of a link in a non-experimental analysis. The ocean or stratospheric dynamics [22], for example, can still render several links in our atmospheric pressure-system study as spurious. For example, the interhemispheric connections from component Nos. 43, 55 seem climatologically implausible. But estimated causal links can yield new hypotheses to be rejected or confirmed by further data analyses involving more variables or model simulations. On the other hand, the finding of non-causality, that is, the absence of a causal link, relies on weaker assumptions. Given that the observed data faithfully represents the underlying process and that potential nonlinearities are powerfully enough captured by the dependence measure, the absence of evidence for a statistical relationship makes it very unlikely that a linking physical mechanism exists. Such findings of non-causality are, therefore, more robust.

Growing data availability promises an unprecedented opportunity for novel insights through causal discovery across all disciplines of science—but causal discovery methods need to overcome some major challenges. Here we focused on linear and nonlinear time-delayed dependencies and addressed high dimensionality and detection biases. We showed that, firstly, our proposed method reaches high detection power and short runtimes making it well suited for large-scale time series datasets. Secondly, the proposed, well-calibrated MCI independence test allows to reliably identify spurious associations and the value of MCI as an estimator of causal strength proves particularly useful to focus on relevant causal dependencies and reduce deteriorating biases. As demonstrated here with an example from the Earth system, these methodological advancements can help to discover novel causal insights also in many other fields of science.

Acknowledgments

We thank B. Hoskins for help with the climatological interpretation and G. Balasis, D. Coumou, J. Donges, F. Fröhlich, J. Haigh, J. Heitzig, M. Kretschmer, J. Kurths, M. Mengel, C.-F. Schleussner, E. van Sebille, and J. Zscheischler for helpful discussions and comments. J.R. is funded by a postdoctoral award by the James S. McDonnell Foundation. The authors thank D. Hartman, J. Hlinka, N. Jajcay, M. Vejmelka, and M. Paluš for providing the climate data components and C. Linstead for help with high-performance computing. We gratefully acknowledge the European Regional Development Fund (ERDF), the German Federal Ministry of Education and Research and the Land Brandenburg for supporting this project by providing resources on the high-performance computer system at the Potsdam Institute for Climate Impact Research. Software is available online under <https://github.com/jakobrunge/tigramite>.

Author contributions

J.R. designed the method, analyzed the data, and prepared the manuscript. D.S. contributed to mathematical formulation. All authors discussed the results and contributed to editing the manuscript.

References

1. D. Bacciu, T. A. Etchells, P. J. G. Lisboa, and J. Whittaker. Efficient identification of independence networks using mutual information. *Comput. Stat.*, 28(2):621–646, 3 2012.
2. L. Barnett, A. B. Barrett, and A. K. Seth. Granger causality and transfer entropy Are equivalent for gaussian variables. *Phys. Rev. Lett.*, 103(23):238701, 12 2009.
3. Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J R Stat Soc Series B Stat Methodol*, 57(1):289–300, 1995.
4. S. Boccaletti, G. Bianconi, R. Criado, C. I. del Genio, J. Gómez-Gardeñes, M. Romance, I. Sendiña-Nadal, Z. Wang, and M. Zanin. The structure and dynamics of multilayer networks. *Phys. Rep.*, 544(1):1–122, 2014.
5. L. Breiman and J. H. Friedman. Estimating Optimal Transformations for Multiple Regression and Correlation. *JASA*, 80(391):580–598, 1985.
6. E. Bullmore and O. Sporns. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.*, 10(3):186–98, 3 2009.
7. T. Choi and M. J. Schervish. On posterior consistency in nonparametric regression problems. *J. Multivar. Anal.*, 98(10):1969–1987, 2007.
8. D. Colombo and M. H. Maathuis. Order-Independent Constraint-Based Causal Structure Learning. *J. Mach. Learn. Res.*, 15:3921–3962, 2014.
9. T. M. Cover and J. A. Thomas. *Elements of Information Theory*. John Wiley & Sons, Hoboken, 2006.
10. J. J. Daudin. Partial association and an measures to qualitative application regression. *Biometrika*, 67(3):581–590, 1980.
11. M. Drouard, G. Rivière, P. Arbogast, M. Drouard, G. Rivière, and P. Arbogast. The Link between the North Pacific Climate Variability and the North Atlantic Oscillation via Downstream Propagation of Synoptic Waves. *J. Climate*, 28(10):3957–3976, 5 2015.
12. I. Ebert-Uphoff and Y. Deng. Causal discovery for climate research using graphical models. *J. Climate*, 25(17):5648–5665, 2012.

-
13. M. Eichler. Graphical modelling of multivariate time series. *Probab. Theory Related Fields*, 153(1):233–268, 2012.
 14. S. R. Flaxman, D. B. Neill, and A. J. Smola. Gaussian Processes for Independence Tests with non-iid Data in Causal Inference. *ACM TIST*, 7(2):22, 2016.
 15. K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel Measures of Conditional Dependence. In *NIPS*, volume 20, pages 489–496, 2007.
 16. C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
 17. A. Gretton, K. Fukumizu, C. H. Teo, L. Song, B. Schölkopf, and A. J. Smola. A kernel statistical test of independence. In *NIPS*, volume 21, pages 585–592, 2008.
 18. A. E. Hoerl and R. W. Kennard. Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*, 12(1):55–67, 1970.
 19. B. J. Hoskins and T. Ambrizzi. Rossby Wave Propagation on a Realistic Longitudinally Varying Flow. *J. Atmos. Sci.*, 50(12):1661–1671, 1993.
 20. B. J. Hoskins and D. J. Karoly. The Steady Linear Response of a Spherical Atmosphere to Thermal and Orographic Forcing. *J. Atmos. Sci.*, 38(6):1179–1196, 1981.
 21. P. Hoyer, D. Janzing, J. Mooij, J. Peters, and B. Schölkopf. Nonlinear causal discovery with additive noise models. In *NIPS*, volume 15, pages 689–696, 2009.
 22. J. W. Hurrell, Y. Kushnir, G. Ottersen, and M. Visbeck. An Overview of the North Atlantic Oscillation. In J. W. Hurrell, Y. Kushnir, G. Ottersen, and M. Visbeck, editors, *The North Atlantic Oscillation: Climatic Significance and Environmental Impact*, pages 1–35. American Geophysical Union, Washington, D. C., 2003.
 23. D. Janzing, D. Balduzzi, M. Grosse-Wentrup, and B. Schölkopf. Quantifying causal influences. *Ann. Stat.*, 41(5):2324–2358, 10 2013.
 24. H. F. Kaiser. The varimax criterion for varimax rotation in factor analysis. *Psychometrika*, 23(3):187–204, 1958.
 25. E. Kalnay, M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu, M. Chelliah, W. Ebisuzaki, W. Higgins, J. Janowiak, K. C. Mo, C. Ropelewski, J. Wang, A. Leetmaa, R. Reynolds, R. Jenne, and D. Joseph. The NCEP/NCAR 40-year reanalysis project. *B. Am. Meteorol. Soc.*, 77(3):437–471, 3 1996.
 26. H. Kantz and T. Schreiber. *Nonlinear Time Series Analysis*. Cambridge University Press, Cambridge, 2003.
 27. D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge, 2010.
 28. A. Kraskov, H. Stögbauer, and P. Grassberger. Estimating mutual information. *Phys. Rev. D*, 69(6):066138, 2004.
 29. M. Kretschmer, D. Coumou, J. F. Donges, and J. Runge. Using causal effect networks to analyze different arctic drivers of midlatitude winter circulation. *J. Climate*, 29(11):4069–4081, 3 2016.
 30. S. L. Lauritzen. *Graphical models*. Oxford University Press, Oxford, 1996.
 31. Y. Li and M. G. Genton. Single-Index Additive Vector Autoregressive Time Series Models. *Scand. J. Stat.*, 36(3):369–388, 9 2009.
-

-
32. D. Lopez-Paz, K. Muandet, B. Schölkopf, and I. Tolstikhin. Towards a Learning Theory of Cause-Effect Inference. In *ICML*, volume 32, 2015.
 33. M. Mader, W. Mader, L. Sommerlade, J. Timmer, and B. Schelter. Block-bootstrapping for noisy data. *J. Neurosci. Methods*, 219(2):285–91, 10 2013.
 34. J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, Cambridge, 2000.
 35. F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.*, 12(Oct):2825–2830, 2011.
 36. M. Peifer, B. Schelter, B. Guschlbauer, B. Hellwig, C. H. Lücking, and J. Timmer. On Studentising and Blocklength Selection for the Bootstrap on Time Series. *Biom. J.*, 47(3):346–357, 6 2005.
 37. J. Peters, D. Janzing, and B. Schölkopf. Causal Inference on Time Series using Restricted Structural Equation Models. In *NIPS*, pages 154–162, 2013.
 38. J. Peters, D. Janzing, and B. Schölkopf. Causal Inference on Time Series using Restricted Structural Equation Models. In *NIPS*, pages 154–162, 2013.
 39. B. Pompe and J. Runge. Momentary information transfer as a coupling measure of time series. *Phys. Rev. E.*, 83(5):051122, 5 2011.
 40. C. Rasmussen and C. Williams. *Gaussian processes for machine learning*. MIT Press, Cambridge, MA, USA, 2006.
 41. D. N. Reshef, Y. A. Reshef, H. K. Finucane, S. R. Grossman, G. McVean, P. J. Turnbaugh, E. S. Lander, M. Mitzenmacher, and P. C. Sabeti. Detecting Novel Associations in Large Data Sets. *Science*, 334(6062):1518–1524, 12 2011.
 42. D. N. Reshef, Y. A. Reshef, P. C. Sabeti, and M. M. Mitzenmacher. An Empirical Study of Leading Measures of Dependence. *arXiv:1505.02214 [stat.ME]*, 5 2015.
 43. J. Runge. Quantifying information transfer and mediation along causal pathways in complex systems. *Phys. Rev. E*, 92(6):062829, 2015.
 44. J. Runge, R. V. Donner, and J. Kurths. Optimal model-free prediction from multivariate time series. *Phys. Rev. E*, 91(5):052909, 2015.
 45. J. Runge, J. Heitzig, N. Marwan, and J. Kurths. Quantifying causal coupling strength: A lag-specific measure for multivariate time series related to transfer entropy. *Phys. Rev. E*, 86(6):061121, 2012.
 46. J. Runge, J. Heitzig, V. Petoukhov, and J. Kurths. Escaping the Curse of Dimensionality in Estimating Multivariate Transfer Entropy. *Phys. Rev. Lett.*, 108(25):258701, 2012.
 47. J. Runge, V. Petoukhov, J. F. Donges, J. Hlinka, N. Jajcay, M. Vejmelka, D. Hartman, N. Marwan, M. Paluš, and J. Kurths. Identifying causal gateways and mediators in complex spatio-temporal systems. *Nat. Comm.*, 6:8502, 2015.
 48. J. Runge, V. Petoukhov, and J. Kurths. Quantifying the Strength and Delay of Climatic Interactions: The Ambiguities of Cross Correlation and a Novel Measure Based on Graphical Models. *J. Climate*, 27(2):720–739, 2014.
 49. T. Schreiber. Measuring information transfer. *Phys. Rev. Lett.*, 85(2):461–464, 2000.
 50. D. Sejdinovic, B. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Ann. Stat.*, 41(5):2263–2291, 10 2013.

-
51. P. Spirtes and C. Glymour. An Algorithm for Fast Recovery of Sparse Causal Graphs. *Soc. Sci. Comput. Rev.*, 9(1):62–72, 1991.
 52. P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*, volume 81. The MIT Press, Boston, 2000.
 53. E. V. Strobl and P. L. Spirtes. Estimating and Controlling the False Discovery Rate for the PC Algorithm Using Edge-Specific P-Values. *arXiv:1607.03975*, (July), 7 2016.
 54. G. Sugihara, R. May, H. Ye, C.-h. Hsieh, E. Deyle, M. Fogarty, and S. Munch. Detecting causality in complex ecosystems. *Science*, 338(6106):496–500, 10 2012.
 55. J. Sun, D. Taylor, and E. M. Bollt. Causal Network Inference by Optimal Causation Entropy. *SIAM J Appl Dyn Syst.*, 14(1):27, 2014.
 56. R. Tibshirani. Regression shrinkage and selection via the lasso. *J R Stat Soc Series B Stat Methodol*, 58(1):267–288, 1996.
 57. A. Tikhonov. Regularization of incorrectly posed problems. *Soviet Mathematics Doklady*, 4(0):1624–1627, 1963.
 58. M. Vejmelka, L. Pokorná, J. Hlinka, D. Hartman, N. Jajcay, and M. Paluš. Non-random correlation structures and dimensionality reduction in multivariate climate data. *Clim. Dynam.*, 44(9-10):2663–2682, 7 2014.
 59. T.-c. Yeh. On Energy Dispersion in the Atmosphere. *J. Met.*, 6(1):1–16, 2 1949.
 60. K. Zhang, J. Peters, D. Janzing, and B. Schölkopf. Kernel-based Conditional Independence Test and Application in Causal Discovery. In *UAI*, pages 804–813, 2011.

Supplementary Material

Detecting causal associations in large nonlinear time series datasets

Jakob Runge, Dino Sejdinovic, and Seth Flaxman

Contents

S1 Causal discovery	13
S1.1 Two-step causal discovery method (PCMCI)	13
S1.2 Alternative methods	15
S1.3 Pre-whitening and block-shuffle autocorrelation remedies	16
S1.4 Conditional independence test implementations (ParCorr and GPACE)	16
S2 Properties of PCMCI	18
S2.1 Definition of causal strength	18
S2.2 MCI measures causal strength	19
S2.3 MCI largely removes detection biases	20
S2.4 MCI yields a well-calibrated conditional independence test	20
S2.5 PCMCI yields an asymptotically consistent estimate of the time series graph	21
S3 Analytical comparison with alternative measures	21
S3.1 General analysis	21
S3.2 Analytical model example setup	23
S3.3 Autocorrelation bias (Fig. S1B)	23
S3.4 Common driver bias (Fig. S1C)	24
S4 Numerical experiments comparing PCMCI with alternative methods	24
S4.1 Model setup	24
S4.2 Analysis of detection biases (Fig. 2 and Supplementary Figs. S6,S7,S8)	26
S4.3 Analysis of high-dimensional networks (Fig. 3 and Supplementary Figs. S9,S10)	27
S4.4 Sample size analysis (Figs. S11,S12)	27
S4.5 Analysis of denser network models with correlated noise terms (Figs. S13,S14)	27
S5 Numerical experiments comparing MCI with alternative measures	28
S5.1 Model setup	28
S5.2 Analysis of alternative autocorrelation remedies (Fig. S15)	29
S5.3 Non-gaussian noise (Fig. S16)	29
S5.4 Stronger and weaker common driver forcing strength (Fig. S17)	29
S5.5 Sample size analysis (Fig. S18)	29
S6 Tigramite software package	30
S7 Earth system example	30
S7.1 Data and analysis setup	30
S7.2 Details for Fig. 4 and Supplementary Figs. S19,S20,S21,S22,S23,S24	31
S8 Algorithms	32
S9 Supplementary Tables	33
S10 Supplementary Figures	38

S1 Causal discovery

S1.1 Two-step causal discovery method (PCMCI)

In our framework, the dependency structure of a set of time series variables is represented in a graphical model [30]. While the type of graph depicted in Fig. 1B is easier to visualize, it does not represent the spatio-temporal dependency structure underlying complex dynamical systems which is more comprehensively grasped in a time series graph [13] as shown in Fig. 1C. If, for example, graphical models are estimated without taking into account lagged variables, associations can easily be confounded by the influence of common drivers at past times.

As shown in Fig. 1C, the nodes of a time series graph are defined as the variables $X^1, X^2, \dots \in \mathbf{X}$ at different times $t, t-1, t-2, \dots$ up to a maximum time lag τ_{\max} and a link $X_{t-\tau}^i \rightarrow X_t^j$ for $\tau > 0$ exists if $X_{t-\tau}^i$ and X_t^j are *not* conditionally independent given the past of the whole process,

$$X_{t-\tau}^i \not\perp\!\!\!\perp X_t^j \mid \mathbf{X}_t^- \setminus \{X_{t-\tau}^i\}, \quad (\text{S1})$$

where ' $\not\perp\!\!\!\perp$ ' denotes the absence of a conditional independence and $\mathbf{X}_t^- \setminus \{X_{t-\tau}^i\} = (\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \dots) \setminus \{X_{t-\tau}^i\}$ the past of the multivariate process excluding $X_{t-\tau}^i$. Assuming stationarity, the links are repeated for every $t' < t$ if a link exists at time t . The parents \mathcal{P} of a variable are defined as the set of all nodes with a link towards it. Contemporaneous dependencies for $\tau = 0$ can be defined in different ways [43]. Here they are left undirected, but recent techniques [32, 37, 52] could be applied to determine causal directionality for contemporaneous links.

Our causal discovery technique to estimate the time series graph is based on a two-step procedure:

1. Obtain an *approximate* estimate of the parents for all variables $X_t^j \in \mathbf{X}_t = (X_t^1, X_t^2, \dots, X_t^N)$ with Algorithm S1
2. Use these approximate parents as conditions in the MCI causal discovery Algorithm S2, which tests all variable pairs $(X_{t-\tau}^i, X_t^j)$ with $i, j \in \{1, \dots, N\}$ and time-delay $0 < \tau \leq \tau_{\max}$ implementing the MCI test

$$\text{MCI: } X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{P}(X_t^j) \setminus \{X_{t-\tau}^i\}, \mathcal{P}_{p_X}(X_{t-\tau}^i), \quad (\text{S2})$$

where $\mathcal{P}_{p_X}(X_{t-\tau}^i) \subseteq \mathcal{P}(X_{t-\tau}^i)$ denotes the p_X strongest parents according to the sorting in Algorithm S1.

One can also restrict the maximum number of parents used for $\mathcal{P}(X_t^j)$, but here we impose no restrictions.

Algorithm S1 in the first step is the skeleton-discovery part of the PC algorithm [51] in its more robust modification called PC-stable [8] and adapted to time series. In summary, for every variable $X_t^j \in \mathbf{X}_t$ the algorithm starts by initializing the preliminary parents $\tilde{\mathcal{P}}(X_t^j) = \mathbf{X}_t^- = (\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \dots, \mathbf{X}_{t-\tau_{\max}})$ and then iteratively removes parents $X_{t-\tau}^i$ from $\tilde{\mathcal{P}}(X_t^j)$ if the null hypothesis

$$\text{PC: } X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{S} \quad (\text{S3})$$

is accepted at a significance threshold α , where \mathcal{S} are different combinations of subsets of $\tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$ with iteratively increasing cardinality (up to a maximum number of combinations q_{\max}). The algorithm converges for a link $X_{t-\tau}^i \rightarrow X_t^j$ once $\mathcal{S} = \tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$ and the null hypothesis $X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$ is rejected. Our fast variant is obtained by restricting the maximum number of combinations q_{\max} per iteration to a very small integer value. The MCI step is based on the information-theoretic measure *momentary information transfer* introduced in [39, 45]).

As listed in Tab. S1, the free parameters of this method (in addition to free parameters of the conditional independence test statistic) are the maximum time delay τ_{\max} , the significance threshold α and the maximum number of combinations q_{\max} per iteration in Algorithm S1, and the maximum number p_X of conditions of the driver variable in Algorithm S2. We abbreviate different parameter choices by $\text{PC}_{q_{\max}}(\alpha) + \text{MCI}_{p_X}$, if not clear from the context. In Figs. S4, S5 we show the overall performance of PCMCI for different choices of q_{\max} , α , and p_X .

Choice of τ_{\max} The maximum time delay depends on the application and should be chosen according to the maximum physical time lag expected in the complex system. We recommend a rather large choice that includes peaks in the lagged cross-correlation function (or a more general measure), because a too large choice of τ_{\max} merely leads to longer runtimes of PCMCI, but not to an increased estimation dimension as for GC.

Choice of q_{\max} We observe hardly any difference between using $q_{\max} = 1$ and $q_{\max} = 2$ in the condition selection algorithm in our numerical experiments. Because in our version of the PC-stable algorithm, we sort the preliminary parents $\tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$ after each iteration, restricting the maximum number of combinations to $q_{\max} = 1$ will result in only the strongest driver(s) being used as a condition \mathcal{S} in Eq. S3 for every dimension. While testing more combinations ($q_{\max} = 2$) can help to slightly bring down false parents among the conditions, this choice results in longer runtimes and we found the effect to be negligible regarding false positive rates (FPRs) of the two-step PCMCI method here.

Choice of α Too small values of α result in many true links not being included in the condition set for the MCI tests and, hence, increase FPRs of PCMCI. For $\alpha > 0.1$ in our numerical experiments, we find that the FPRs stabilize around the expected level. Too high levels of α lead to high dimensionality of the condition set which reduces detection power and strongly increases the runtime. For denser causal networks than in the experiments shown in Figs. 2,3, we found that a larger value of α helps to more reliably control FPRs (Figs. S13,S14), which, however, takes longer and leads to a larger number of parents increasing the estimation dimension. Note that for a very liberal threshold α in Algorithm S1, essentially all $N\tau_{\max}$ variables would be selected as conditions and the MCI test becomes a GC test.

Choice of p_X While the parents $\mathcal{P}(X_t^j)$ are important to assess conditional independence, the additional conditions $\mathcal{P}_{p_X}(X_{t-\tau}^i) \subseteq \mathcal{P}(X_{t-\tau}^i)$ are necessary to account for autocorrelation and the causal strength scaling as analyzed in Sect. S2. To limit high dimensionality, we strongly restrict the number of conditions $\mathcal{P}_{p_X}(X_{t-\tau}^i)$ with the free parameter p_X . In our tests we found that a small value $p_X = 1$ or $p_X = 3$ already suffices to reduce inflated FPRs due to strong autocorrelation and estimate causal strength, resulting in a confined scaling of power with link strength (see, e.g., Figs. S6,S7). The reason is that typically the largest driver will be the autodependency and conditioning out its influence already diminishes the effect of strong autocorrelations.

In general, a good choice of q_{\max} , α , and p_X depends on how sensitive the conditional independence test is to slightly larger estimation dimensions. Since the ParCorr and GPACE (see, e.g., Fig. S5) statistics are not that sensitive in a range way below the high dimensionality of the GC test, we choose $q_{\max} = 1$, $p_X = 1$ or $p_X = 3$, and a not too small α to capture more true parents. For a test statistic like conditional mutual information [46], the estimation dimension can be more strongly limited by a smaller value of α . Then one can also restrict the maximum condition dimension p_{\max} , that is, the maximum cardinality of \mathcal{S} , in Algorithm S1 and also the maximum number of parents in $\mathcal{P}(X_t^j)$ for the MCI tests in Algorithm S2.

In [55] an alternative algorithm to reconstruct the parents based on forward-selection and backward elimination is discussed. In general, any algorithm that provides a consistent estimate of the causal parents can be used.

False discovery rate control The PC algorithm can also be combined with *false discovery rate* [3] controls [1], while here we use the FDR pruning only for the p-values resulting from the MCI step for the whole

Table S1. Parameters of PCMCI method.

Parameter	Description
q_{\max}	Max. number of combinations in Algorithm S1
α	Significance threshold in Algorithm S1
p_X	Max. number of parents of $X_{t-\tau}^i$ in Algorithm S2
τ_{\max}	Max. time delay in Algorithms S1,S2

time series graph (see real data examples). The false discovery approach controls the expected proportion of discoveries (rejected null hypotheses) that were false.

Computational complexity The computational complexity of the first step of PCMCi (Algorithm S1) strongly depends on the network structure, but it can be limited by adapting the parameters α , q_{\max} , or also the maximum condition dimension p_{\max} . The sparser the causal dependencies, the faster the convergence. Our version of the PC algorithm is speed up with $q_{\max} = 1$. The MCI step (Algorithm S2) involves $N^2 \tau_{\max}$ tests (for $\tau > 0$) of a maximal dimensionality of $2 + |\mathcal{P}(X_t^j)| + p_X$. The computational time will then depend on how the conditional independence test scales with this dimensionality and on the sample length T . In Figs. S9,S10,S11,S12,S13,S14 we analyze runtimes of the different techniques for different network sizes N and sample sizes T . GC is comparably slow for large network sizes because of the scaling with high dimensionality, while the original PC algorithm is slow because of a large number of conditional independence tests since many different combinations of sets of conditions are tested.

Our method is freely available in the **Tigramite** python software package from <https://github.com/jakobrunge/tigramite>. Section S6 provides more details.

S1.2 Alternative methods

Granger causality In its original formulation, Granger causality between time series variables X and Y is based on fitting a linear or nonlinear model including possible confounders to Y and a causal link $X \rightarrow Y$ is assessed by quantifying whether the inclusion of the past of variable X in the model significantly reduces the prediction error about Y [16]. Here, we use a general, lag-specific version based directly on the definition of time series graphs, that is, by conditioning the independence test between nodes $X_{t-\tau}^i$ and X_t^j on the whole past of the process:

$$\text{GC: } X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathbf{X}_t^- \setminus \{X_{t-\tau}^i\}, \quad (\text{S4})$$

where $\mathbf{X} = (X^1, X^2, \dots)$ denotes the multivariate process and $\mathbf{X}_t^- = (\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \dots)$ its past. In practice, the past is truncated at a maximum time lag τ_{\max} .

PC algorithm The original PC algorithm was formulated for general random variables without assuming a time order. It consists of several phases where first, in the skeleton-discovery phase, an undirected graphical model [30] is estimated whose links are then oriented using a set of rules [51,52]. We implement the skeleton-discovery phase of the PC algorithm as given in Algorithm S1, but in contrast to our fast variant, the original version does not restrict the number of condition combinations to test. Here we use a large choice of $q_{\max} = 10$. In contrast to GC or PCMCi, it is not straightforward to assess the confidence of causal links with the PC algorithm because links are iteratively removed. Here we use a p-value assessment for causal links according to [53],

$$p(X_{t-\tau}^i \rightarrow X_t^j) = \max_{\{\mathcal{S}\}} p(X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{S}), \quad (\text{S5})$$

that is, the maximum of all p-values from the conditional independence tests for different condition sets \mathcal{S} in Eq. S3 defines the aggregated p-value of a causal link. This causal discovery method we term PC in the numerical experiments. In the Supplement, we also show results without pruning links based on p-values, but where all links that ‘survived’ the algorithm are counted as significant (PC(raw)). We found that the p-value pruning tends to be over-conservative and FPR levels cannot be reliably controlled just below a desired threshold.

MCI₀ In the MCI step of the PCMCi method, one can also test

$$\text{MCI}_0: X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{P}(X_t^j) \setminus \{X_{t-\tau}^i\} \quad (\text{S6})$$

for all links, that is, the MCI test without the condition on the parents $\mathcal{P}(X_{t-\tau}^i)$, or, equivalently, $p_X = 0$. Note that for $X_{t-\tau}^i \in \mathcal{P}(X_t^j)$ this is the test in the last step of the PC algorithm, but here also links that were removed in Algorithm S1 are tested again. For weak autocorrelation, we found that this variant provides more power, especially for nonlinear dependencies, which is likely due to the conditioning set having lower dimensionality. However, common driver forcings can still lead to considerable detection biases (see Sect. S3 for analytical examples).

For conditional mutual information as a test statistic this measure was called *information transfer to Y* (ITY) in [45]. For a link $X \rightarrow Y$, it quantifies the unique information in X entering a process Y without excluding the information in the past of X . ITY and the information-theoretic version of MCI called *momentary information transfer* (MIT) are part of a whole suite of measures to quantify causal information flow in complex systems [43]. Another measure called ITX quantifies the unique information emanating from a process as an information-theoretic analog to Sims causality, and several other measures quantify mediation on information pathways as discussed in [43].

Bivariate MCI In the example applications, we also evaluate a bivariate version of the MCI test,

$$\text{bivMCI: } X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{P}_{\text{auto}}(X_t^j) \setminus \{X_{t-\tau}^i\}, \mathcal{P}_{\text{auto}}(X_{t-\tau}^i). \quad (\text{S7})$$

where $\mathcal{P}_{\text{auto}}(X_t^j) = (\mathcal{P}(X_t^j) \setminus \{X_{t-\tau}^i\}) \cap \{X_{t-1}^j, X_{t-2}^j, \dots\}$ denotes only those parents belonging to the past of X_t^j and correspondingly for $X_{t-\tau}^i$. Thus, this test removes autocorrelation to some extent, but does not exclude common drivers or indirect dependencies. Also autodependencies induced by common drivers are not removed.

Unconditional pairwise measures In the real data examples, we also investigate the unconditional pairwise measures $I(X_{t-\tau}; Y_t)$. In the ParCorr implementation this is simply the Pearson correlation coefficient and in the GPACE implementation the Maximal Correlation coefficient as estimated by ACE [5].

Comparing MCI causal networks and link values with Correlation, bivMCI, and MCI₀ to some extent allows to assess whether a non-significant link in the MCI test is due to an indirect dependency or a common driver, on the one hand, or more due to autocorrelations.

S1.3 Pre-whitening and block-shuffle autocorrelation remedies

We also tested the MCI₀ test Eq. (S6) using a pre-whitening (PCpw) and a block-shuffle permutation test (PCbs). For the pre-whitening test, we preprocessed all N time series by estimating the univariate lag-1 autocorrelation coefficients $\hat{a}_i = \rho(X_{t-1}^i; X_t^i)$ and regressing out the AR(1) autocorrelation part of the signals:

$$\tilde{X}_t^i = X_t^i - \hat{a}_i X_{t-1}^i \quad \forall t \text{ and } i = 1, \dots, N. \quad (\text{S8})$$

Then the PC₁+MCI₀ test is applied to these residuals $\tilde{\mathbf{X}}$.

Another remedy is a computationally expensive block-shuffle permutation test instead of the tabulated critical values. We use a block-shuffle surrogate test here following [36] and [33]. For the test statistic I , an ensemble of $M = 500$ values of $I(X_{t-\tau}^*; Y_t \mid \dots)$ is generated where $X_{t-\tau}^*$ is a block-shuffled surrogate of $X_{t-\tau}$, i.e., with blocks of the original time series permuted. As an optimal block-length we use the method described in [36] and [33] for non-overlapping blocks. The optimal block-length formula Eq. (6) in [33] involves the decay rate of the envelope of the autocorrelation function $\gamma(\tau)$. The latter was estimated up to a maximum delay of 5% of the samples, and the envelope was estimated using the Hilbert transform. Then a function $C\phi^\tau$ was fit to the envelope with constant C to obtain the decay rate ϕ . The block length was limited to a maximum of 10% of the sample length. Finally, the estimated values are sorted, and a p-value is obtained as the fraction of surrogates with values greater or equal than the estimated value.

S1.4 Conditional independence test implementations (ParCorr and GPACE)

Similarly to the PC algorithm, the framework (Algorithm S1 and S2) that we propose can be used in conjunction with any conditional independence test – these will typically be based on estimating different

dependence measures with associated test statistics. Here we implement two tests: partial correlation (ParCorr) applicable to the multivariate Gaussian case which can only capture linear dependence and a novel nonlinear two-step conditional independence test we term GPACE.

ParCorr Partial correlation for testing $X \perp\!\!\!\perp Y \mid Z$ is implemented with ordinary least squares (OLS) regression of X and Y on Z assuming the model

$$\begin{aligned} X &= h_X(Z) + \epsilon^X \\ Y &= h_Y(Z) + \epsilon^Y, \end{aligned} \tag{S9}$$

for linear functions h and independent Gaussian noise terms ϵ . All variables are standardized before the regression (mean subtracted and divided by standard deviation). Then the partial correlation test is given by testing the Pearson correlation of the residuals with a t -test with $T - D_Z - 2$ degrees of freedom where D_Z is the dimensionality of Z . In Fig. S9 we also discuss ridge regression [18, 57] and Lasso [56] estimators implemented in the `scikit-learn` python library ([35], version 0.18). Ridge regression with leave-one-out-cross-validated regularization parameter $\alpha \in \{0.1, 1, 2, \dots, 500\}$ was applied on standardized time series (no intercept fitted). Then the degrees of freedom are estimated as $T - \sum_{k=1}^{D_Z} \lambda_k^2 / (\lambda_k^2 + \alpha)$ where λ_k are the singular values of Z . For Lasso regression a multi-task lasso technique was used. The regularization parameter α was chosen from $\{0.0001, 0.001, 0.01, 0.1, 1\}$ using 5-fold cross-validation. Otherwise, default parameters were used, except for additionally normalizing the data to make the hyper-parameters more robust and restricting a maximum number of iterations to 100 to limit computational time, which was by far the highest of all methods. The degrees of freedom for the subsequent t -test are then given by the sample size T minus the number of non-zero coefficients (same for regressions on X and Y in this technique).

GPACE GPACE is similar to [21], where such framework is termed RESIT (Regression with Subsequent Independence Test). Instead of a linear regression in Eq. S9, the first step is a Gaussian process (GP) regression [40] and the second step a test for the (unconditional) independence of the regression residuals using an estimator of the Maximal Correlation coefficient [5]. GP regression is a widely used Bayesian nonparametric regression approach. Its frequentist counterpart is called kernel ridge regression (KRR) and sometimes also referred to as least squares support vector machine (LS-SVM). Gaussian process regression was implemented using `sklearn` with Radial Basis Function (RBF) kernel and an added nugget parameter 0.1. The bandwidth was estimated using maximum likelihood estimation with the `fmin_cobyla` optimizer. However, while [21] and a similar approach in [14] both use tests based on Hilbert-Schmidt Independence Criterion (HSIC) [17] for the second step, we propose to use the nonparametric alternating conditional expectation (ACE) estimator of the maximal correlation [5] based on copula-transformed residual time series yielding uniform marginal distributions. The copula transformation allows to pre-compute the distribution (implemented in `Tigramite`) and critical values (see Tab. S4) under the independence null hypotheses – thereby avoiding computationally expensive permutation tests needed for HSIC or conditional mutual information estimates [28]. ACE also has linear computational complexity in the time series sample length (in contrast to the quadratic cost of HSIC and conditional mutual information), and was found to have good statistical power in [42]. Unlike HSIC, where an additional kernel choice on the residuals needs to be specified, ACE does not require further tuning parameters. GP regression, unfortunately, scales $\sim T^3$ with the sample length making it computationally prohibitive for large sample sizes. Then sparse GP models can be used [40].

Assumptions Any two-step procedure for conditional independence testing, which includes a regression followed by an unconditional test on the regression residuals, has an underlying assumption of a functional dependence on the conditioning variables as given by Eq. S9. In the presence of dependencies which cannot be represented even in a nonlinear functional form h , regression will not be able to remove these dependencies on the conditioning variables. Then the two-step procedure should be replaced with alternative techniques to measuring and testing conditional dependence, e.g., kernel conditional independence tests [15, 60] or estimates of conditional mutual information based on nearest-neighbor mutual information estimators [28]. These approaches, however, come with complex estimation procedures of the null distribution, typically do not scale well to high dimensionality of the conditioning variable and have a number of additional tuning parameters.

Multivariate extensions We note that we focus on the case of univariate time series X and Y as is typical in causal discovery applications (while the conditioning variable Z can and mostly will be multivariate). However, our method can be readily extended to multivariate time series X and Y , as well as those taking values in structured or non-Euclidean domains if suitable conditional independence tests are available. In this case one could indeed use a two-step procedure involving vector-valued regression and HSIC (since ACE is no longer applicable) or the above-mentioned kernel conditional independence tests [15, 50, 60]. In complex dynamical systems, multivariate variables can, for example, result from delay-embeddings to better unfold the dynamics of nonlinear systems [26]. Or, one can create multivariate variables from aggregating time series of multiple subprocesses to represent a dynamical process on a higher layer. Causal networks reconstructed on such different layers then lend themselves to analysis with measures from the growing field of multilayer networks [4].

S2 Properties of PCMCI

In this section, we provide more formal definitions and proofs for the properties of PCMCI stated in the main article.

S2.1 Definition of causal strength

Our definition of causal strength is based on an additive time series model here. Denote

$$\begin{aligned} X_{t-\tau} &= g_X(\mathcal{P}(X_{t-\tau})) + \eta_{t-\tau}^X \\ Y_t &= g_Y(\mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) + \tilde{\eta}_t^Y, \end{aligned} \quad (\text{S10})$$

where g_X and g_Y are deterministic functions and the parents $\mathcal{P}(X_{t-\tau})$ and $\mathcal{P}(Y_t)$ can include external variables as well as autodependencies, for example, $Y_{t-1} \in \mathcal{P}(Y_t)$ and $X_{t-\tau-1} \in \mathcal{P}(X_{t-\tau})$. Note that $\mathcal{P}(Y_t)$ can additionally depend on X at other time delays than $t - \tau$. $\mathcal{P}(X_{t-\tau})$ and $\mathcal{P}(Y_t)$ may also overlap.

Suppose the “innovation terms” satisfy the following assumptions:

- **(A1)** $\{\tilde{\eta}_t^Y, \eta_{t-\tau}^X\}$ are jointly independent of $\mathcal{P}(Y_t) \setminus \{X_{t-\tau}, \mathcal{P}(X_{t-\tau})\}$ given $\mathcal{P}(X_{t-\tau})$,
- **(A2)** $\eta_{t-\tau}^X \perp\!\!\!\perp \mathcal{P}(X_{t-\tau})$
- **(A3)** $\tilde{\eta}_t^Y \perp\!\!\!\perp \{\mathcal{P}(X_{t-\tau}), \eta_{t-\tau}^X\} | X_{t-\tau}$

Now a possible link $X_{t-\tau} \rightarrow Y_t$ in this model can be represented by including a dependence on $X_{t-\tau}$ in $\tilde{\eta}_t^Y$:

- Linear additive dependency: $\tilde{\eta}_t^Y = cX_{t-\tau} + \eta_t^Y$
- Nonlinear additive dependency: $\tilde{\eta}_t^Y = f(X_{t-\tau}) + \eta_t^Y$ for some nonlinear deterministic function f
- General nonlinear dependency: $\tilde{\eta}_t^Y = f(X_{t-\tau}, \eta_t^Y)$

Here η_t^Y denotes an independent noise term for which we assume

- **(A4)** $\eta_t^Y \perp\!\!\!\perp \{\mathcal{P}(Y_t), \mathcal{P}(X_{t-\tau})\}$.

(A2) and (A4) imply that both η_t^Y and $\eta_{t-\tau}^X$ are independent in time. (A3) implies that the term $\tilde{\eta}_t^Y$ need not be independent of $\mathcal{P}(X_{t-\tau})$. However, it can depend on $\mathcal{P}(X_{t-\tau})$ and $\eta_{t-\tau}^X$ only through $X_{t-\tau}$. Note that $\tau > 0$ so that there is no ambiguity in the direction of the link $X_{t-\tau} \rightarrow Y_t$, if present. The assumption in (A1) that $\eta_{t-\tau}^X \perp\!\!\!\perp \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\} | \mathcal{P}(X_{t-\tau})$ implies that there can be no dependence of $\eta_{t-\tau}^X$ with other parents of Y , which is further discussed in [43].

We now investigate an information-theoretic definition of causal strength based on conditional mutual information [9]. Consider the following mutual information decomposition of X with the dependence term $\tilde{\eta}$ in Y :

$$I(X_{t-\tau}; \tilde{\eta}_t^Y) = I(\mathcal{P}(X_{t-\tau}), \eta_{t-\tau}^X; \tilde{\eta}_t^Y) \quad (\text{S11})$$

$$= I(\mathcal{P}(X_{t-\tau}); \tilde{\eta}_t^Y) + \underbrace{I(\eta_{t-\tau}^X; \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau}))}_{\text{momentary causal strength}} \quad (\text{S12})$$

Proof The first line follows from the Markov chain $\{\mathcal{P}(X_{t-\tau}), \eta_{t-\tau}^X\} \rightarrow X_{t-\tau} \rightarrow \tilde{\eta}_t^Y$ and hence $H(\tilde{\eta}_t^Y | X_{t-\tau}) = H(\tilde{\eta}_t^Y | X_{t-\tau}, \mathcal{P}(X_{t-\tau}), \eta_{t-\tau}^X)$ and the fact that $X_{t-\tau}$ is a deterministic function of $\{\mathcal{P}(X_{t-\tau}), \eta_{t-\tau}^X\}$. The second equality follows by the chain rule of (conditional) mutual information [9]. ■

We will focus on the second term $I(\eta_{t-\tau}^X; \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau}))$, which we define as the (*momentary*) *causal strength*, since it directly corresponds to “momentary” dependence in $\tilde{\eta}_t^Y$ on $X_{t-\tau}$ and does not come through the parents of $X_{t-\tau}$. In the case of absence of the link $X_{t-\tau} \rightarrow Y_t$, $I(X_{t-\tau}; \tilde{\eta}_t^Y) = 0$ so the causal strength $I(\eta_{t-\tau}^X; \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau}))$ is also zero. There are several proposals for measures of causal strength, see, for example, [23]. Our definition of causal strength is based on the fundamental concept of *source entropy* as further discussed in [43]. In Sect. S3 we compare momentary causal strength with GC and other measures analytically.

S2.2 MCI measures causal strength

We first investigate MCI in its most general information-theoretic version and then discuss under which assumptions our results also hold for ParCorr and GPACE. Information-theoretically, the MCI test can be implemented with conditional mutual information and the test statistic is then called *momentary information transfer* [39, 43, 45] (MIT). For model (S10) it can be analyzed as follows:

$$I_{X \rightarrow Y}^{\text{MCI}}(\tau) = I(X_{t-\tau}; Y_t | \mathcal{P}(X_{t-\tau}), \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (\text{S13})$$

$$= I(g_X(\mathcal{P}(X_{t-\tau})) + \eta_{t-\tau}^X; g_Y(\mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) + \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau}), \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (\text{S14})$$

$$= I(\eta_{t-\tau}^X; \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau}), \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (\text{S15})$$

$$= I(\eta_{t-\tau}^X; \tilde{\eta}_t^Y | \mathcal{P}(X_{t-\tau})), \quad (\text{S16})$$

where Eq. S15 follows from translational invariance of CMI [9] and Eq. S16 from (A1). That is, for processes of the class given in model (S10), MCI is an estimator of our notion of causal strength. ■

For the dependence types above, MCI can be further simplified:

- Linear additive dependency:

$$I_{X \rightarrow Y}^{\text{MCI}}(\tau) = I(\eta_{t-\tau}^X; c\eta_{t-\tau}^X + \eta_t^Y) \quad (\text{S17})$$

since by (A2, A4) $\mathcal{P}(X_{t-\tau})$ is independent of $\eta_{t-\tau}^X, \eta_t^Y$.

- Nonlinear additive dependency:

$$I_{X \rightarrow Y}^{\text{MCI}}(\tau) = I(\eta_{t-\tau}^X; f(X_{t-\tau}) + \eta_t^Y | \mathcal{P}(X_{t-\tau})) \quad (\text{S18})$$

- General nonlinear dependency:

$$I_{X \rightarrow Y}^{\text{MCI}}(\tau) = I(\eta_{t-\tau}^X; f(X_{t-\tau}, \eta_t^Y) | \mathcal{P}(X_{t-\tau})) \quad (\text{S19})$$

Thus, for the linear additive dependency, where causal strength can be attributed to a single coefficient c , MCI depends only on this coefficient and on the noise terms, but not on g_Y, g_X . MCI is then independent of biases due to the parents $\mathcal{P}(X_{t-\tau})$ and $\mathcal{P}(Y_t)$, which could include autodependencies. A causal signal can, thus, be better detected against noise coming from confounding drivers or autocorrelation as demonstrated in analytical examples in Sect. S3. On the other hand, for the nonlinear cases, there can still be various dependencies because the function f mixes $\eta_{t-\tau}^X$ with $\mathcal{P}(X_{t-\tau})$.

Further assumptions for ParCorr and GPACE The richness of the classes of functions g_X, g_Y and f allowed in the model S10 will depend on the specific approach used in conditional independence testing within the MCI framework. For ParCorr, g_X, g_Y and f all must be linear and the noise terms Gaussian. On the other hand, for GPACE, g_X and g_Y can be nonlinear but sufficiently smooth functions. Namely, when testing for conditional independence between $X_{t-\tau}$ and Y_t given $Z_t = \mathcal{P}(X_{t-\tau}) \cup \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}$ using GPACE, we postulate $X_{t-\tau} = h_X(Z_t) + \epsilon_{t-\tau}^X$ and $Y_t = h_Y(Z_t) + \epsilon_t^Y$, where h_X and h_Y are deterministic functions and $\epsilon_{t-\tau}^X$

and ϵ_t^Y are zero-mean random variables independent of Z_t . Note that h_Y does not consist only of g_Y above but it is crucially allowed to also depend on $\mathcal{P}(X_{t-\tau})$ (but not on $X_{t-\tau}$), i.e., $h_Y(Z_t) = g_Y(\mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) + c\tilde{g}_{f,g_X}(\mathcal{P}(X_{t-\tau}))$, for some potentially nonlinear function \tilde{g} of $\mathcal{P}(X_{t-\tau})$. Under this assumption, testing for $X_{t-\tau} \perp\!\!\!\perp Y_t | Z_t$ is equivalent to testing for $\epsilon_{t-\tau}^X \perp\!\!\!\perp \epsilon_t^Y$ [10, 14]. Thus, the analysis extends to GPACE *provided that* the nonlinear functions h_X and h_Y have been unveiled to a sufficient fidelity by the first step in the procedure, i.e., by Gaussian process (GP) regression [40]. The consistency properties of GP regression for specific classes of functions h_X and h_Y have been studied in [7]. A full analysis of GPACE would require considering those learning theoretic guarantees on regression functions and how they impact the properties of the subsequent independence tests (such as ACE), which is beyond the scope of this work.

While nonlinearity of g_X and g_Y is dealt with straightforwardly using nonparametric regression techniques and does not affect the derivation of Eq. S17, for nonlinear functions $f(x)$ such as in Eq. (S18, S19), the noise $\eta_{t-\tau}^X$ cannot be ‘isolated’ anymore from the parents $\mathcal{P}(X_{t-\tau})$ and MCI may also depend on other coefficients, though typically to a lesser extent. This can be seen in the nonlinear numerical examples where the power of MCI is not constant anymore for varying autocorrelation and dimensionality (Fig. S15). Also for linear f , but very large autocorrelations and nonlinearities g_X, g_Y (Fig. S15), strongly non-Gaussian noise (Fig. S16), or large forcings (Fig. S17), the effect cannot be fully regressed out anymore using the present Gaussian process framework. Possibly, a different kernel or adjusting the bandwidth optimization can help in this case for GPACE, or alternative conditional independence tests can be used [15, 28, 60].

S2.3 MCI largely removes detection biases

Next to the sample size and the chosen significance level, statistical power crucially depends on the actual effect size relative to the amount of noise in the data. As proven above, the MCI statistic is an estimator of our notion of causal strength which defines the effect size of interest in our method. Hence, the statistical power of the MCI test statistic is solely related to the causal strength.

For a linear coupling function f in model S10, the causal strength can be attributed to a single coupling coefficient and the detection power of MCI is unbiased regarding other coefficients such as related to autocorrelation or other forcings. This implies that MCI has equal power for causal couplings of equal strength and scales with causal strength. In large-scale causal discovery studies, this important feature implies that links of high causal strength are more likely to be detected, no matter how strongly autocorrelated the involved variables are, or how both are forced by external common drivers, implying an improved causal signal-to-noise ratio. This property is compared with GC and the PC test in analytical examples in Sect. S3.

S2.4 MCI yields a well-calibrated conditional independence test

In the absence of a link $X_{t-\tau} \rightarrow Y_t$ for model (S10), we have $\tilde{\eta}_t^Y = \eta_t^Y$, that is, $\tilde{\eta}_t^Y \perp\!\!\!\perp X_{t-\tau}$. Then MCI is given by

$$I_{X \rightarrow Y}^{\text{MCI}}(\tau) = I(X_{t-\tau}; Y_t | \mathcal{P}(X_{t-\tau}), \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (\text{S20})$$

$$= I(\eta_{t-\tau}^X; \eta_t^Y) \quad (\text{S21})$$

$$= 0. \quad (\text{S22})$$

by (A1-A4). Importantly, the innovation terms $\eta_{t-\tau}^X, \eta_t^Y$ are independent in time. If both are also identically distributed for every t , respectively, the dependence of MCI only on these innovation terms implies that statistical tests on $I_{X \rightarrow Y}^{\text{MCI}}(\tau) = 0$ can be conducted under the *iid-assumption*, where *iid* stands for independent and identically distributed. In the linear case with partial correlation as an independence measure, *iid*-ness allows utilizing the t -distribution with $T - D_Z - 2$ degrees of freedom, while for GPACE the pre-computed sampling distribution from uniform independent samples is justified. This property leads to a well-calibrated test and explains our numerical finding that the MCI test has the expected rate of FPRs even for strongly autocorrelated variables. This natural approach to avoiding time-dependence in the sample we found to be superior to adapting the null-distribution by block-shuffle techniques or pre-whitening as analyzed in Sect. S5. Note, however, that this result is derived here for the population version of MCI and its application to

empirical estimators should be considered with some caution and would rely on consistency and unbiasedness of these estimators.

S2.5 PCMCI yields an asymptotically consistent estimate of the time series graph

The PC algorithm consistently estimates the parents of each variable in the limit of large sample sizes under the assumptions of faithfulness and causal sufficiency [52]. Our scalable approximate variant PC_1 may result in false parents (since not all conditions are exhaustively tested), but still captures all true parents in the limit of large sample sizes under the faithfulness assumption. The false parents would then be detected as spurious in the MCI step. Therefore, PCMCI yields an asymptotically consistent estimate of the time series graph under faithfulness and causal sufficiency. To some extent the assumption of faithfulness can be relaxed for our two-step method because links that are removed in the PC algorithm step are tested again in the MCI step, but this topic is beyond the scope of this paper. As noted above, for a very liberal threshold α in Algorithm S1, essentially all variables would be selected as conditions, and the MCI test becomes a GC test.

S3 Analytical comparison with alternative measures

S3.1 General analysis

In the following, we investigate the dependency structure of alternative measures for model (S10) for the most general information-theoretic case.

Mutual information (MI) Unconditional mutual information, defined as $I_{X \rightarrow Y}^{MI}(\tau) = I(X_{t-\tau}; Y_t)$, trivially depends on g_Y, g_X and, thus, possibly on common parents, i.e., confounders, since no variables are conditioned out. Also, when testing for $X_{t-\tau} \perp\!\!\!\perp Y_t$, one cannot generally assume independent samples as the observations have a temporal dependence, i.e., time itself is a confounder.

Information transfer (ITY) If we exclude the condition on the parents of X as in MCI_0 , we arrive at the information-theoretic measure *information transfer to Y* (ITY) as introduced in [45]. It can be decomposed as

$$I_{X \rightarrow Y}^{ITY}(\tau) = I(X_{t-\tau}; Y_t | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (S23)$$

$$= I(\eta_{t-\tau}^X, \mathcal{P}(X_{t-\tau}); Y_t | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) \quad (S24)$$

$$= I(\mathcal{P}(X_{t-\tau}); Y_t | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) + I(\eta_{t-\tau}^X; Y_t | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}, \mathcal{P}(X_{t-\tau})) \quad (S25)$$

The second term is MCI and, thus, measures causal strength, while the first term explicitly depends on the dependencies between Y_t and $\mathcal{P}(X_{t-\tau})$. Note that from this line it also follows that (for population quantities) $I_{X \rightarrow Y}^{MCI}(\tau) \leq I_{X \rightarrow Y}^{ITY}(\tau)$. If no link exists, we have $Y_t = \eta_t^Y + g_Y(\mathcal{P}(Y_t) \setminus \{X_{t-\tau}\})$ and

$$I_{X \rightarrow Y}^{ITY}(\tau) = I(\mathcal{P}(X_{t-\tau}); \eta_t^Y | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}) + I(\eta_{t-\tau}^X; \eta_t^Y | \mathcal{P}(Y_t) \setminus \{X_{t-\tau}\}, \mathcal{P}(X_{t-\tau})) \quad (S26)$$

$$= 0, \quad (S27)$$

since η_t^Y is assumed independent of $\mathcal{P}(X_{t-\tau})$ and $\eta_{t-\tau}^X$. But because $\mathcal{P}(X_{t-\tau})$ could include non-*iid* variables, a statistical test based on ITY, called MCI_0 in our framework, does not generally fulfill the assumption of independent samples. This explains the increased FPRs obtained in our numerical experiments (e.g., Fig. S15 and in another numerical experiment in Sect. S5).

Granger causality The information-theoretic analogue of Granger causality is called *transfer entropy* [49]. For Gaussian variables they can be shown to be equivalent [2]. We use a lag-specific definition of Granger




Figure S1. Comparison of information-theoretic dependence measures for simple Gaussian model example S34. (A) Time series graph, where the blue boxes denote the condition used in the ITY measure and the red box the additional condition for MCI. MI is unconditioned and GC is conditioned on the whole past up to $\tau_{\max} = 3$. Analytical expressions are given in the text. All innovation terms have unit variance. (B) Dependence of different measures for increasing autocorrelation in X . Here $a = 0$, $c_{WX} = 0$, $c_{XY} = 0.6$, $c_{WY} = 0.3$, $c_{XZ} = 0.3$. Only MCI fully excludes autocorrelation dependencies, MI and ITY are increasing while GC is decreasing for larger autocorrelation in X , implying a lower causal signal-to-noise ratio. Note that MI, ITY, and GC in general also depend on other forcing coefficients. (C) Dependence of different measures for varying the common driver forcing of W . Here $a = 0.3$, $c_{WX} = 0.6$, $c_{XY} = 0.6$, $c_{XZ} = 0.3$. For negative c_{WY} , here MI and GC even tend to zero implying that a true causal link is not detected due to a common driver forcing of opposite sign. Also ITY varies with the forcing, while only MCI has the same constant causal strength value for any forcing.

causality defined in Sect. S1 with $\mathbf{X}_t^- = (\mathbf{X}_{t-1}, \mathbf{X}_{t-2}, \dots)$ as

$$I_{X \rightarrow Y}^{\text{GC}}(\tau) = I(X_{t-\tau}; Y_t | \mathbf{X}_t^- \setminus \{X_{t-\tau}\}) \quad (\text{S28})$$

$$= I(\eta_{t-\tau}^X, \mathcal{P}(X_{t-\tau}); Y_t | \mathbf{X}_t^- \setminus \{X_{t-\tau}\}) \quad (\text{S29})$$

$$= \underbrace{I(\mathcal{P}(X_{t-\tau}); Y_t | \mathbf{X}_t^- \setminus \{X_{t-\tau}\})}_{=0} + I(\eta_{t-\tau}^X; Y_t | \mathbf{X}_t^- \setminus \{X_{t-\tau}\}) \quad (\text{S30})$$

$$= I(\eta_{t-\tau}^X; Y_t | \mathbf{X}_t^- \setminus \{X_{t-\tau}\}) \quad (\text{S31})$$

where the first term vanishes because $\mathcal{P}(X_{t-\tau}) \subset \mathbf{X}_t^- \setminus \{X_{t-\tau}\}$. Here biases arise due to the dependencies of \mathbf{X}_t^- on $\eta_{t-\tau}^X$.

If no link exists, we have $Y_t = \eta_t^Y + g_Y(\mathcal{P}(Y_t) \setminus \{X_{t-\tau}\})$ and

$$I_{X \rightarrow Y}^{\text{GC}}(\tau) = I(\eta_{t-\tau}^X; \eta_t^Y | \mathbf{X}_t^- \setminus \{X_{t-\tau}\}) \quad (\text{S32})$$

$$= 0, \quad (\text{S33})$$

which implies that the GC test is essentially performed on independent samples. Also this result is confirmed in our numerical experiments where the GC test gave expected FPRs also for strongly autocorrelated variables in the ParCorr implementation. However, in the GPACE implementation, we found inflated FPRs for strongly autocorrelated links (Fig. S10), which is likely due to high dimensionality where the autocorrelations are not properly regressed out.

S3.2 Analytical model example setup

To compare measures of link strength, we investigate the link $X_{t-2} \rightarrow Y_t$ in a model of a multivariate linear Gaussian time series (time series graph shown in Fig. S1A):

$$\begin{aligned} Z_t &= aZ_{t-1} + c_{XZ}X_{t-1} + \eta_t^Z \\ X_t &= a_X X_{t-1} + c_{WX}W_{t-1} + \eta_t^X \\ Y_t &= aY_{t-1} + c_{XY}X_{t-2} + c_{WY}W_{t-3} + \eta_t^Y \\ W_t &= aW_{t-1} + \eta_t^W \end{aligned} \quad (\text{S34})$$

with independent Gaussian white noise processes η_i with variances σ^2 . The parents are $\mathcal{P}_{Y_t} = \{Y_{t-1}, X_{t-2}, W_{t-3}\}$ and $\mathcal{P}_{X_{t-2}} = \{X_{t-3}, W_{t-3}\}$. Generally, the conditional mutual information $I(X; Y | \mathbf{Z})$ between two normally distributed random vectors X and Y conditional on a (possibly multivariate) random vector \mathbf{Z} is given by

$$I(X; Y | \mathbf{Z}) = \frac{1}{2} \ln \left(\frac{|\Gamma_{Y\mathbf{Z}}| |\Gamma_{X\mathbf{Z}}|}{|\Gamma_{\mathbf{Z}}| |\Gamma_{YX\mathbf{Z}}|} \right), \quad (\text{S35})$$

where, e.g., $|\Gamma_{Y\mathbf{Z}}|$ is the determinant of the covariance matrix of (Y, \mathbf{Z}) . In Fig. S1 we consider two cases: Autocorrelation bias and a common driver forcing.

S3.3 Autocorrelation bias (Fig. S1B)

For $a = 0$, $c_{WX} = 0$, $c_{XY} = 0.6$, $c_{WY} = 0.3$, $c_{XZ} = 0.3$, the model can be treated analytically (see [45]). Then mutual information is given by

$$I_{X \rightarrow Y}^{\text{MI}} = \frac{1}{2} \ln \left(1 + \frac{(c_{XY}^2 \sigma_X^2) / (1 - a_X^2)}{c_{WY}^2 \sigma_W^2 + \sigma_Y^2} \right), \quad (\text{S36})$$

and ITY is given by

$$I_{X \rightarrow Y}^{\text{ITY}} = \frac{1}{2} \ln \left(1 + \frac{(c_{XY}^2 \sigma_X^2) / (1 - a_X^2)}{\sigma_Y^2} \right). \quad (\text{S37})$$

Thus, MI depends on the coefficients and variances of the other input processes to Y , while ITY (or MCI_0) still depends at least on the autocorrelation coefficient a_X .

The link-specific Granger causality is here given by

$$I_{X \rightarrow Y}^{\text{GC}} = \frac{1}{2} \ln \left(1 + \frac{c_{XY}^2 \sigma_X^2 \sigma_Z^2}{\sigma_Y^2 (c_{XZ}^2 \sigma_X^2 + (1 - a_X^2) \sigma_Z^2)} \right), \quad (\text{S38})$$

which depends on autocorrelation and, counter-intuitively, on how X drives another process Z , that is not a driver of Y .

Finally, MCI is

$$I_{X \rightarrow Y}^{\text{MCI}} = \frac{1}{2} \ln \left(1 + \frac{c_{XY}^2 \sigma_X^2}{\sigma_Y^2} \right), \quad (\text{S39})$$

which corresponds to our definition of causal strength and solely depends on the model coefficients that govern the coupling, i.e., the variances σ_X^2 , σ_Y^2 , and the coupling coefficient c_{XY} .

In Fig. S1 we show the dependence of MI, ITY, GC, and MCI plotted against autocorrelations in X . While MI and ITY are positively biased, GC is negatively biased for increasing autocorrelation. Since statistical power scales with effect size, these dependencies imply that higher autocorrelated links are more likely detected with ITY (or MCI_0) and less likely detected with GC for this example. This can partially explain the lower power in GC for higher autocorrelated links observed in Fig. 2 and is an example of what we call a low causal signal-to-noise ratio. Note that this bias is in addition to the low power of GC due to high dimensionality.

S3.4 Common driver bias (Fig. S1C)

For the coefficients $a = 0.3$, $c_{WX} = 0.6$, $c_{XY} = 0.6$, $c_{XZ} = 0.3$, we investigate how the common driver W affects the different measures if it induces a forcing of opposite sign between X_{t-2} and Y_t with $c_{WY} < 0$ while the causal link coefficient c_{XY} is positive. This example was evaluated numerically with very long sample sizes ($T = 200,000$).

For this case MI (as would a linear correlation) and GC even tend to zero implying that a true causal link is not detected due to a low causal signal-to-noise ratio. Also ITY varies with the forcing, while only MCI has the same constant causal strength value for any forcing as expected from our analytical results (Sect. S2). In our application examples, we observe several cases of low correlation values with a higher MCI value, which can at least partially be explained by this analytical finding.

Note that the case of a vanishing MI (here for large negative c_{WY}) constitutes an example of a violation of the faithfulness assumption made in the PC algorithm [52]. The link $X_{t-2} \rightarrow Y_t$ would be removed in Algorithm S1 in the first step and X_{t-2} would be missing in the parents of Y_t . While this pathological case can introduce spurious associations and biases in PCMCI, the link $X_{t-2} \rightarrow Y_t$ would likely be detected with PCMCI, since all links are tested again in the MCI step.

S4 Numerical experiments comparing PCMCI with alternative methods

S4.1 Model setup

To compare the different causal discovery methods, we use the following model: For $i, j \in \{1, \dots, N\}$ we randomly choose L pairs (i, j) with $i \neq j$ and generate time series according to

$$X_t^j = a_j X_{t-1}^j + \sum_i c_i f_i(X_{t-\tau_i}^i) + \eta_t^j \quad (\text{S40})$$

for $j = 1, \dots, N$ and with

- a_j uniformly randomly drawn from $\{0, 0.2, 0.4, 0.6, 0.8, 0.9\}$ for one half of the ensemble and from $\{0.6, 0.8, 0.9, 0.95\}$ for the other more autocorrelated half
- *iid* noise $\eta^j \sim \mathcal{N}(0, 1)$
- $f_i(x)$ uniformly drawn from smooth real linear functions $f^{(1)}(x) = x$ with 50% probability (100% for ParCorr experiments) and nonlinear functions $f^{(2)}(x) = (1 - 4e^{-x^2/2})x$ and $f^{(3)}(x) = (1 - 4x^3e^{-x^2/2})x$, each with 25% probability for the GPACE experiments
- τ_i uniformly randomly drawn from $\{1, 2\}$
- c_i such that in the corresponding single-link model with $X_t^i = \eta_t^i$ and $X_t^j = cf(X_{t-1}^i) + \eta_t^j$ we have a constant mutual information $I(\eta_{t-1}^j; cf(\eta_{t-1}^i) + \eta_t^j) = I_{X \rightarrow Y}$, with the sign of c_i positive or negative with equal probability; Tab. S3 lists the corresponding coefficients

Here $I_{X \rightarrow Y}$ measures the link strength in *nats* and is set constant for all L links in a given model. Note that it corresponds to our notion of causal strength, as defined in Sect. S2, only for linear dependencies (coupling function $f^{(1)}$). For nonlinear functions, our notion of causal strength depends also on the parents of the driver variable, while the link strength with corresponding coefficient given in Tab. S3 is defined in the corresponding counterfactual single-link model without any drivers. Since the model suffices the assumptions stated in Sect. S2, our notion of causal strength for any of the coupling functions defined above can be estimated by MCI. To guarantee stationarity, the functions $f_i(x)$ are all linear in the limit of large x and we dismiss all sets of coefficients for which the corresponding vector autoregressive model with nonlinear functions f_i replaced by linear ones is nonstationary according to a unit root test [31]. Fig. S2 gives an example realization for nonlinear models with $N = 5$.

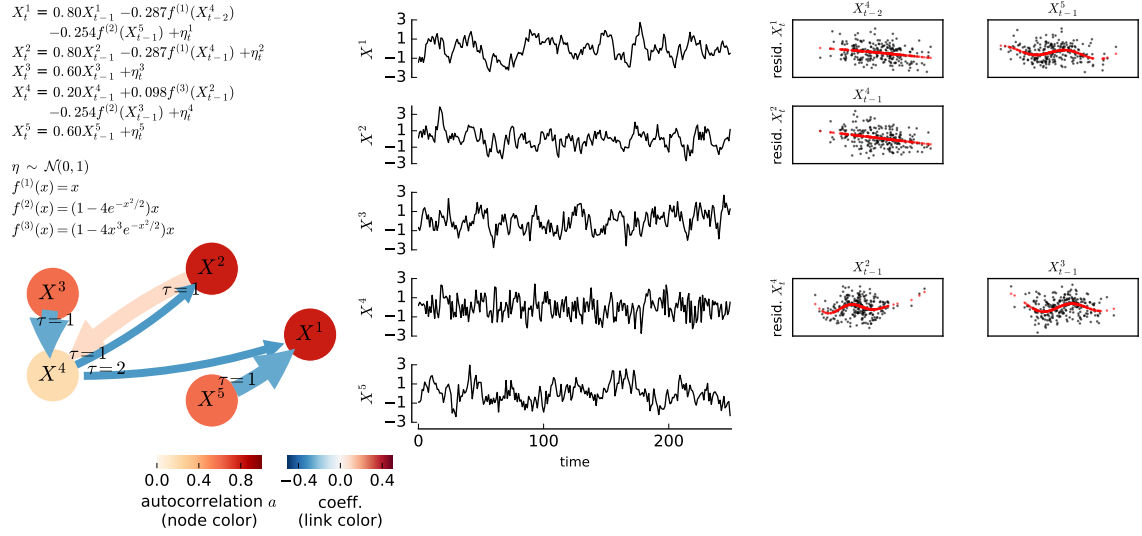


Figure S2. Example realization for $N = 5$ with $L = N$ links between different variables and $T = 250$ for the nonlinear network model described in Sect. S4. The equations are given in the upper left, with independent zero-mean unit variance Gaussian noise terms $\eta \sim \mathcal{N}(0, 1)$. In the graph, thick arrows correspond to nonlinear links. Even though the links have different colors corresponding to the coupling coefficient (see Tab. S3), their link strength described in Sect. S4 is the same for all links. The node color depicts the value of the autocorrelation coefficients. The scatter plots are taken between the residual $X_t^j \mid \mathcal{P}(X_t^j) \setminus \{X_{t-\tau}^i\}$ (after a Gaussian Process regression of X_t^j on its remaining parents) and $X_{t-\tau}^i$. The red line shows the GP regression prediction as examples of the coupling functions $f^{(1)}$, $f^{(2)}$, $f^{(3)}$ occurring in the equations on the left.

S4.2 Analysis of detection biases (Fig. 2 and Supplementary Figs. S6,S7,S8)

For the numerical experiments shown in Fig. 2 and Figs. S6,S7,S8, for N variables a total of 120 models with different coupling topologies of $L = N$ links of different type, time lag, and link strength were generated: For ParCorr we study only linear couplings $cf^{(1)}(x) = cx$, where c is adjusted to obtain the desired link strengths $I_{X \rightarrow Y} = 0.02, \dots, 0.08$ *nats* (see Tab. S3). The total number of links in this numerical experiment for each link strength is $40 \cdot L$. For GPACE we investigate three different coupling types: Linear $cf^{(1)}(x) = cx$ as before (half of the links for each $I_{X \rightarrow Y}$) and nonlinear $cf^{(2)}(x) = c(1 - 4e^{-x^2/2})x$ and $cf^{(3)}(x) = c(1 - 4x^3e^{-x^2/2})x$ (a quarter of the links each). For all methods, the maximum lag was set to $\tau_{\max} = 5$ leading to a total number of false cross-links of $\tau_{\max}N(N-1) - L$ per model. Additionally, the individual variables have different autocorrelations as given above. With $L = N$ cross-links in each model, we have an average degree of 1 (plus a possible autodependency). For ParCorr we set $N = 20$ and for GPACE $N = 5$. The graph in Fig. 2A depicts an example network for $N = 20$ and a link strength of $I_{X \rightarrow Y} = 0.08$ with the black arrow width proportional to the true positive rate (TPR) and the gray arrows to the false positive rates (FPRs). In this graph FPRs are averaged across all lags).

To assess TPRs and FPRs for the individual links in each model, 100 realizations with sample length $T = 150$ for ParCorr and $T = 250$ for GPACE were generated (if not stated otherwise). Note that the error in the estimate of a FPR of 0.05 (or a TPR of 0.95) is roughly $\sqrt{0.05(1-0.05)/100} \approx 0.02$. In the graph plot in Fig. 2A, the significance level for PCMCI was adjusted to obtain the same average FPR as the PC algorithm in the graph above. The bottom rows in Fig. 2B,C show boxplots of the distribution of FPRs and the upper rows of the TPR for linear and nonlinear dependencies where $I_{X \rightarrow Y}$ is increasing along the x-axis in each plot. Only cross-links were considered here. The left and right boxplot depict the distributions for all weakly autocorrelated pairs with mean autodependency $(I(X_{t-1}; X_t) + I(Y_{t-1}; Y_t))/2 < 0.7$ among the two variables X and Y of a link, and for strongly autocorrelated pairs $((I(X_{t-1}; X_t) + I(Y_{t-1}; Y_t))/2 \geq 0.7)$, respectively. The boxes show the 25-75% and whiskers the 1-99% percentile range, the median is marked by a bar and the mean with ‘x’. Note the logarithmic y-axis in the bottom panel for $\text{FPR} > 0.1$.

In the ParCorr implementation (Fig. S6), the high-dimensional GC leads to very low power levels. The original PC-stable algorithm for $q_{\max} = 10$ and $\alpha = 0.2$ is shown here with pruned links where the p-value for each link is assessed as defined in [53] (PC) and without pruning links where all links that ‘survived’ the algorithm are counted as significant [PC(raw)]. Both variants show heterogeneous and mostly lower TPRs even for strong causal links. Clearly, no confined scaling of TPRs with link strength is visible. GC and PC are compared with PCMCI for different numbers of additional MCI conditions: $\text{PC}_1 + \text{MCI}$ for $p_X = 1$ and $\text{PC}_1 + \text{MCI}_3$ for $p_X = 3$. Here including more conditions in $\mathcal{P}(X_{t-\tau}^j)$ with $\text{PC}_1 + \text{MCI}_3$ clearly leads to lower FPRs and also slightly improves the scaling of TPRs with link strength. Here we also show a combination of the PC_1 -preselection step with a subsequent MCI_0 test as described in Eq. (S6) that excludes the condition on the parents of X . As expected from the results shown in another numerical experiment in Fig. S15, the FPR for MCI_0 is higher, especially for stronger autocorrelated links. It is, therefore, not surprising that the TPRs are the highest throughout. $\text{PC}_1 + \text{MCI}_0$ shows some scaling with link strength, but not as confined as for MCI shown to the right. The last panel depicts $\text{PC}_1 + \text{MCI}_0$ applied on pre-whitened (pw) time series. Even though only AR(1) autocorrelations are present, pre-whitening does not significantly reduce FPRs. There are even partially higher FPRs and the detection power is much lower. This result is further investigated in Fig. S15.

In the GPACE implementation (Fig. S7), GC shows a surprising resistance to high dimensionality with FPRs only slightly higher than for MCI. But its power is lower, and also the scaling with link strength is not as confined. The original PC-stable algorithm also here has difficulties to control FPR levels and has very heterogeneous power. Here higher values of p_X are not as efficient for $\text{PC}_1 + \text{MCI}$ regarding FPRs. $p_X = 1$ suffices in our analysis to control FPRs and obtain a scaling of power levels with link strength for linear links even for strongly autocorrelated pairs.

As discussed in Sect. S2, the power-scaling feature does not hold for more nonlinear links, here strongly nonlinear links of coupling type $cf^{(3)}$ (top row) are not well detected with MCI or the GC test. Interestingly, for the coupling type $cf^{(3)}$ we observe that stronger autocorrelated links are on average less well detected than weaker autocorrelated pairs, while for the coupling type $cf^{(2)}$ it is the opposite (even more evident in Fig. S12). Different types of nonlinearities, thus, can have very different effects on detection power biases. For larger sample size, detection rates for nonlinear links are much better (Fig. S8).

PC₁+MCI₀ has the highest FPRs here, but it is also the only one that has high TPRs for the strongly nonlinear coupling type $f^{(3)}$. For weakly autocorrelated variables the FPR is the same as for MCI and MCI₀ is, thus, preferable to achieve higher power, which is likely due to the conditioning set having lower dimensionality. However, common driver forcings can still lead to considerable detection biases (see Sect. S3 for analytical examples).

S4.3 Analysis of high-dimensional networks (Fig. 3 and Supplementary Figs. S9,S10)

For the numerical experiments comparing high-dimensional performance for $N = 5, 10, 20, 40, 60$ shown in Figs. 3,S9,S10,S13,S14, we investigate 10 models for a link strength $I_{X \rightarrow Y} = 0.04$.

With $L = N$ links in each model, we have an average degree of 1 for all network sizes (plus an autodependency). To assess TPRs and FPRs for the individual links in each model, 100 realizations with sample length $T = 150$ for ParCorr and $T = 250$ for GPACE were generated. The tick labels on the top note the average runtime and its standard deviation across the different model setups. The runtime estimates were evaluated on Intel Xeon E5-2667 v3 8C processors with 3.2GHz. The runtimes of the different methods are not excessively increasing with network size due to an efficient residual recycling implementation in **Tigramite**, where already computed (Gaussian process) residuals are saved in memory to be re-used.

As shown in Fig. S9, GC clearly fails with low power levels already for $N = 5$ and cannot be evaluated anymore for $N \geq 40$, where the estimation dimension with $\tau_{\max} = 5$ is about 200 and larger than $T = 150$ samples. We also depict results for the regularized GC techniques via ridge and lasso regression. Both of these give higher power than the unregularized regression, but the FPRs are over-controlled and still the power levels steadily decrease much stronger than for PCMCI, especially for ridge regression. The results for the original PC algorithm show that FPRs are hard to control with its threshold parameter α . PC(raw) refers to the PC algorithm without pruning of links. PC₁+MCI avoids high estimation dimensions and is remarkably stable even for very large network sizes with only a slight decrease in power and well-controlled FPRs which can further be reduced by taking into account more of the conditions in MCI₃ with $p_X = 3$.

In the GPACE implementation (Fig. S10), Gaussian process regression demonstrates its remarkable immunity to high dimensions: Even at a dimensionality of $N\tau_{\max} = 60 \cdot 5 = 300$ with just $T = 250$ samples, GC has only slightly higher FPRs and less power than for $N = 5$. Nevertheless, GC still features higher FPRs, especially for high autocorrelation links, and lower power compared to PCMCI. A further advantage of PCMCI, next to the power-scaling property, is its much shorter runtime. GC takes over $7h$ for $N = 60$. The original PC algorithm has an average runtime of $2.1h$ for this network size, which is only slightly higher than that of PCMCI, but the standard deviation of the runtime is $3.5h$ compared to $0.9h$ for PCMCI showing that for some network topologies the PC algorithm converges extremely slowly.

S4.4 Sample size analysis (Figs. S11,S12)

To investigate results for different sample sizes, we ran the same model setup as for the experiments in Fig. S6 for ParCorr and in Fig. S7 for GPACE, but with the link strength fixed to $I_{X \rightarrow Y} = 0.02$ and for varying T . Here 10 different coupling topologies were tested to limit computational time since GPACE scales $\sim T^3$.

For ParCorr, power levels go up with larger samples (Fig. S11) as expected for all methods. Still, PCMCI always outperforms GC and PC. FPRs in all methods except for PC are better controlled for larger sample sizes. For PC(raw), that is, without pruning p-values, the FPRs even increase for larger sample sizes. For larger sample sizes also nonlinear links (Fig. S12) are better detected with PCMCI as expected.

S4.5 Analysis of denser network models with correlated noise terms (Figs. S13,S14)

In Figs. S13,S14 we show a model with $L = 2N$, i.e., an average degree of 2 and where the noise terms η in the model described in Sect. S4 are spatially correlated with a covariance of $\sigma_{ij} = 0.2$ for $i, j \in \{1, \dots, N\}$. That is, the time series graph has also contemporaneous links. The link strength was fixed to $I_{X \rightarrow Y} = 0.02$ and $T = 300$. Here we found that a larger significance level α and a higher p_X better control FPRs.

S5 Numerical experiments comparing MCI with alternative measures




Figure S3. Model for testing the influence of autocorrelation and high dimensionality on independence tests. The model parameters are described in Sect. S5 and Tab. S3. In short, the autocorrelation a (same for all variables) and the number of common drivers D_Z are varied, while the coefficients b and σ_Z are chosen such that the unconditional dependence stays the same, and we only investigate the effect of increasing serial correlation and higher dimensionality. In the GC test the estimation dimension is $1 + \tau_{\max}(D_Z + 2)$, here for $\tau_{\max} = 5$. The conditions of the PC test are given by the gray and blue boxes with dimensionality $D_Z + 3$. The red box denotes the additional condition used in the MCI test with dimensionality $D_Z + 4$. The link labels denote the coefficients. In the model setup all links correspond to linear dependencies, except for the setup shown in the top row for GPACE, where the coupling link with coefficient c is nonlinear $cf^{(2)}(x) = c(1 - 4e^{-x^2/2})x$.

To investigate the isolated effect of autocorrelation and high dimensionality on the GC, PC, and MCI conditional independence tests, we here provide results for an additional numerical experiment. We also show results for alternative autocorrelation remedies as described in Sect. S1.

S5.1 Model setup

The model time series graph is depicted in Fig. S3. In the model setup all links correspond to linear dependencies, except for the setup shown in the top row in the GPACE figures, where the coupling link with coefficient c is nonlinear $cf^{(2)}(x) = c(1 - 4e^{-x^2/2})x$ (functions shown in Fig. S2). The autocorrelation a is the same for all variables $X, Y, Z_1, \dots, Z_{D_Z}$. The coupling coefficient c of the link $X_{t-1} \rightarrow Y_t$ is zero to test FPRs and nonzero to test power as measured by TPRs. For nonzero c its value is chosen such that in the corresponding bivariate model without autocorrelation ($D_Z = 0, a = 0$) we obtain a constant mutual information $I(X_{t-1}; Y_t) = 0.03$ nats for the linear coupling with ParCorr and $I(X_{t-1}; Y_t) = 0.04$ nats for GPACE corresponding to a constant link strength for the coupling function $f^{(1)}$. For the nonlinear coupling $f^{(2)}$, we set $I(X_{t-1}; Y_t) = 0.08$ nats. On the other hand, the common driver forcing coefficient $b = b(D_Z, a)$ and the covariance among the drivers $\sigma_Z = b/2$ are chosen such that in the full model for every pair ($D_Z > 0, a \geq 0$) with $c = 0$ we have $I(X_{t-1}; Y_t) = 0.4$ nats, a relatively strong forcing corresponding to a correlation of ≈ 0.7 . Table S3 lists the coefficients corresponding to different link strengths and forcings. In Fig. S17 we also investigate weaker and stronger common driver forcings.

This setup guarantees that the unconditional dependence stays the same and we only investigate the effect of increasing autocorrelation and higher dimensionality. We test additive Gaussian noise terms $\eta \sim \mathcal{N}(0, 1)$ except for Fig. S16, where we also show results for extreme-value noise terms distributed according to a one-parameter Weibull distribution $\eta \sim (-\ln(U))^{2/3}$ where U is uniformly distributed on $(0, 1]$. For this distribution, we set the link strength to 0.05 nats. Note that the link strength as chosen here corresponds to our

notion of causal strength only in the linear case. For $D_Z = 0$ the model setup corresponds to two autocorrelated processes without a common driver forcing. The sample lengths are $T = 150$ for ParCorr and $T = 250$ for GPACE. 1,000 realizations were evaluated to assess false and true positives.

S5.2 Analysis of alternative autocorrelation remedies (Fig. S15)

In comparing the MCI test with the PC test (Eq. S3) used in the PC algorithm and with Granger causality (GC), we find in Fig. S15 that the MCI and GC tests yield reliable FPRs up to high dimensions while the PC test displays inflated FPRs for increasing autocorrelation, also for $D_Z = 0$ without common drivers. For the same reason also the power of the PC tests increases with autocorrelation, while—at least for linear dependencies—MCI has unbiased power levels. The finding that the PC algorithm ends up with lower power for higher autocorrelated links (e.g., Fig. 2) can be explained as follows: Each test $X_{t-\tau}^i \perp\!\!\!\perp X_t^j \mid \mathcal{S}$ in the PC algorithm (Algorithm S1) has a higher than expected FPR for strongly autocorrelated variable pairs. Therefore, such a link is tested with a larger number of different conditions than weakly autocorrelated pairs, for which the algorithm converges faster. Now a larger number of tests on average leads to a higher p-value of the link, since in the method of [53] the maximum p-value among all tests is chosen. Higher p-values then explain the lower power. For this example also GC is not affected by autocorrelation, but as analyzed in the analytical example in Sect. S3, the GC statistic can also depend on autocorrelation. The GC test has lower power in higher dimensions for ParCorr while GPACE is less affected. For $D_Z = 32$ the dimensionality of the GC test is 171 and exceeds the time series length for ParCorr. While the ParCorr GC test then cannot be evaluated anymore with traditional estimators, GPACE shows still some performance in such high dimensions.

In the top plot in Fig. S15, we depict the results for a nonlinear coupling. In this case, the power of MCI does not stay at the same level for all autocorrelations, but also increases. As discussed further in Sect. S2, the reason is that the autocorrelation contribution then cannot be isolated anymore. This feature of MCI only works well for linear dependencies while still allowing for nonlinear common driver forcing.

The alternative autocorrelation remedies such as a block-shuffle permutation test and pre-whitening do not work well. The problem here is that we test conditional independence in a very general multivariate setting with multiple drivers which are themselves autocorrelated. Pre-whitening or block-shuffling, we find, does not account for such a general non-*iid* structure.

S5.3 Non-gaussian noise (Fig. S16)

In Fig. S16, we investigate the same model with Weibull noise terms. We find robust results for FPRs except for very high dimensions which are avoided anyways with PCMCI. Here the power-scaling property of MCI is more fragile especially for very high autocorrelations where GPACE cannot fully regress out autocorrelations since it assumes Gaussian noise.

S5.4 Stronger and weaker common driver forcing strength (Fig. S17)

In Figs. S17, we test different forcing strengths I of the common drivers Z , that is, different coefficients b (see Tab. S3). While weaker forcing improves results, for strong forcing and very high autocorrelations, GPACE cannot fully regress out autocorrelations and higher FPRs and non-constant power levels for MCI emerge. Here more adapted choices of kernel functions might help. Still, MCI has much lower FPRs than the PC test and slightly less than GC in high dimensions. For $D_Z = 32$ and strong forcing all methods break down but note that PCMCI avoids such a high dimensionality in contrast to GC.

S5.5 Sample size analysis (Fig. S18)

Last, in Fig. S18 we find that the unbiasedness of MCI is even more robust for larger sample sizes, as expected.

Table S2. Sample selection choices in Tigramite.

Masking choice	$I(X_{t-\tau}; Y_t \{Z_{t-\tau_i}^i\})$ estimated only for
X	$t-\tau \in \mathcal{T}_X$
Y	$t \in \mathcal{T}_Y$
Z	$t-\tau_i \in \mathcal{T}_{Z^i} \forall i$
X, Y	$t-\tau \in \mathcal{T}_X \wedge t \in \mathcal{T}_Y$
X, Z^i	$t-\tau \in \mathcal{T}_X \wedge t-\tau_i \in \mathcal{T}_{Z^i} \forall i$
Y, Z^i	$t \in \mathcal{T}_Y \wedge t-\tau_i \in \mathcal{T}_{Z^i} \forall i$
X, Y, Z^i	$t-\tau \in \mathcal{T}_X \wedge t \in \mathcal{T}_Y \wedge t-\tau_i \in \mathcal{T}_{Z^i} \forall i$

S6 Tigramite software package

In the following section, we explain how the real data example was analyzed using the **Tigramite** software package (current version 3.0). **Tigramite** is a time series analysis python module for linear and nonlinear causal inference available from <https://github.com/jakobrunge/tigramite>. **Tigramite** comes with a module that contains several plotting functions to generate high-quality plots of time series, lag functions, causal graphs as depicted in Fig. 1B and time series graphs as in Fig. 1C. All such figures in this article were plotted with this module. Documentation can be found on the repository site. In the following, we explain the important issue of masked samples and detail the real data examples.

Sample selection / masking Time series often contain missing values, or one wishes to only analyze selected samples, for example, only winter months in climate data analyses. **Tigramite** allows to flexibly constrain the analysis on selected samples. For a general test statistic measure $I(X; Y | \{Z^i\})$ for possibly more than one condition Z^i , each argument X, Y, Z^i can be constrained to selected time samples only. For the set of all time indices \mathcal{T} we denote the subsets of selected indices as $\mathcal{T}_X, \mathcal{T}_Y, \mathcal{T}_{Z^i} \subseteq \mathcal{T}$ (which can also comprise the whole set). Then one can estimate $I(X_{t-\tau}; Y_t | \{Z_{t-\tau_i}^i\})$ only from selected time indices t as described in Tab. S2. For example, the choice to mask only Y can be used to measure the influence of any lagged sample in $X_{t-\tau}$ on Y_t at selected samples only, for example, only on winter months in a climate time series. Then $X_{t-\tau}$ can still contain samples from summer months for larger lags τ .

If values are missing, one possible approach is to discard only masked samples in all variables X, Y, Z^i . This will use as many available samples as possible. It may, however, introduce a selection bias in the following way: Suppose X and Y have no missing values, while Z contains missing values, and we obtain an estimate of $I(X; Y)$. If we now consider the estimation of $I(X; Y | Z)$, we not only have added a condition Z , but we also use only a subset of samples in X and Y , which may cause a selection bias. Then one cannot know whether the difference between $I(X; Y)$ and $I(X; Y | Z)$ comes from the condition Z or the subset of samples used. To circumvent this issue, one can include in the mask of Y all samples t in $\mathcal{T}_X, \mathcal{T}_Y, \mathcal{T}_Z$ if a value is missing in *any* of them. Additionally, we have to exclude all t' for $t < t' \leq t + \tau_{\max}$ in $\mathcal{T}_X, \mathcal{T}_Y, \mathcal{T}_Z$. This restricts $I(X; Y)$ to the same set of samples as $I(X; Y | Z)$.

S7 Earth system example

S7.1 Data and analysis setup

The climatological surface-pressure reanalysis dataset [25] studied here can be downloaded from <http://www.esrl.noaa.gov>. The dimension-reduction technique is comprehensively described in [58]. In short, it is based on *Varimax-rotated principal components* [24] and a subsequent significance test to eliminate components merely representing noise. Code for the dimension reduction step is available at <https://github.com/vejmelkam/ndw-climate/blob/master/scripts>. The loadings and time series of all components are shown in the Supplementary Information in [47].

The dataset then consists of $N = 60$ component time series. We choose a time resolution of $\Delta t = 3$ days to study synoptic weather-scale interactions and choose a maximum time delay $\tau_{\max} = 7$ since atmospheric

processes typically decay within few weeks. This choice reflects a balance between resolving causal directionality and averaging over too noisy observations on daily or hourly time scales. In addition, a too high time resolution constitutes a multiple testing challenge if too many lags are considered.

Climatic processes are of a highly seasonal nature implying an at least periodic stationarity of the underlying causal interactions. We, therefore, focus on the winter season here by considering only samples that fall within December to February using the **Tigramite** masking choice Y (see Tab. S2). Taking into account only the more reliable satellite-period, we end up with $T = 969$ samples for 1981–2012. We run PCMCI with $q_{\max} = 1$ and $\alpha = 0.4$ here since we observed a denser interaction structure and found that a higher α better controls FPRs in this case, at the expense of a higher dimensionality (Fig. S14). As an independence test, we apply the ParCorr implementation focusing on linear dependencies with much greater power than with GPACE. To analyze drivers of the NAO we also tested nonlinear links with GPACE. The significance tests were run with $p_X = 3$ at a 5% false discovery (FDR) level using the Benjamini-Hochberg correction [3]. The pre-selection step gave on average $|\mathcal{P}(X_t^j)| = 48$ parents (minimum 36, maximum 58) for use as a condition in the second MCI step, which implies an average estimation dimensionality in the MCI step of about $2 + p_X + |\mathcal{P}(X_t^j)| = 53$ that is much below the dimensionality of GC of $N\tau_{\max} = 420$.

S7.2 Details for Fig. 4 and Supplementary Figs. S19,S20,S21,S22,S23,S24

We compare the MCI analysis with Pearson correlation, bivMCI, GC, and MCI_0 . In Fig. S19, we depict scatter plots for MCI against these measures focusing on lagged links for $\tau > 0$. Correlation and the MCI tests correspond to very different underlying independence hypotheses. Correlation links can differ from MCI links due to inflated autocorrelation, common drivers, or indirect dependencies. These explain the large differences between both measures. In comparing MCI with the bivariate MCI values, at least autocorrelation is partially excluded and the difference between bivMCI and MCI stems mostly from common drivers or indirect links. The MCI vs. MCI_0 scatter plot is naturally more similar since the only difference lies in the exclusion of autocorrelation in the driving variable. Nevertheless, for individual links, this can lead to strong differences in line with our numerical experiments. While MCI and GC correspond to the same underlying conditional independence hypothesis, their test statistic values have a different interpretation as analyzed in Sect. S3 and GC has much lower power explaining the differences shown in this scatter plot. Figure S21 shows that the MCI results are robust also for using another α in the condition-selection algorithm. In Fig. S22 we depict the component regions for the strongest links. In Tab. S6 we provide p-values and FDR-adjusted q-values for the strongest lagged links.

In Fig. S20 we show the corresponding scatter plots for contemporaneous links at lag $\tau = 0$, which are left undirected here. Note that FDR adjustment was conducted only among all contemporaneous links. A large number of links are contemporaneous here, which can possibly be causally resolved on a finer, but noisier, time-resolution or by a further analysis step with recently developed methods that address the inference of causality if no time-ordering is given at all [32, 38, 52]. Since we do not assess causality among contemporaneous interactions, and only excluded lagged effects, many of these links are likely spurious correlations.

In Fig. S23, we show some more results for the analysis of the North Atlantic component No. 5 where p-values are given in Tab. S5. We also studied drivers of the NAO within the same analysis setup but in the GPACE implementation (Fig. S24). Then we also find the link from No. 2 and the probably spurious interhemispheric links from Nos. 43,55, but not the links from Nos. 44,45. These may be due to the lower power of GPACE, or, because a nonlinear common driver induced these links.

S8 Algorithms

Algorithm S1 Pseudo-code for condition-selection algorithm based on the skeleton-discovery step of PC-stable to estimate parents of X_t^j ; with $p_{\max} = N\tau_{\max}$ (i.e., no restriction on the maximum number of parents) and a small integer value for q_{\max} we use this algorithm as a pre-selection step in PCMCI; for the original PC-stable, we set q_{\max} to a large value of 10.

Require: Time series dataset $\mathbf{X} = (X^1, X^2, \dots, X^N)$, selected variable X^j , maximum time lag τ_{\max} , significance threshold α , maximum condition dimension p_{\max} , maximum number of combinations q_{\max} , conditional independence test function

```

1: function CI( $X, Y, Z$ )
2:   Test  $X \perp\!\!\!\perp Y \mid Z$  using test statistic measure  $I$ 
3:   return p-value, test statistic value  $I$ 
4: Initialize preliminary set of parents  $\tilde{\mathcal{P}}(X_t^j) = \{X_{t-\tau}^i : i = 1, \dots, N, \tau = 1, \dots, \tau_{\max}\}$ 
5: Initialize dictionary of test statistic values  $I^{\min}(X_{t-\tau}^i \rightarrow X_t^j) = \infty \ \forall X_{t-\tau}^i \in \tilde{\mathcal{P}}(X_t^j)$ 
6:  $p = -1$ 
7: while  $p < p_{\max}$  and  $|\tilde{\mathcal{P}}(X_t^j)| > p + 1$  do
8:    $p = p + 1$ 
9:   for all  $X_{t-\tau}^i$  in  $\tilde{\mathcal{P}}(X_t^j)$  do
10:     $q = 0$ 
11:    while  $q < q_{\max}$  and still untested conditions  $\mathcal{S} \subseteq \tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$  of dimension  $p$  for the link  $X_{t-\tau}^i \rightarrow X_t^j$  exist do
12:       $q = q + 1$ 
13:      Choose new condition  $\mathcal{S}$  of cardinality  $p$  in lexicographic order from the sorted parents  $\tilde{\mathcal{P}}(X_t^j) \setminus \{X_{t-\tau}^i\}$ 
14:      Run CI test to obtain (p-value,  $I$ )  $\leftarrow$  CI( $X_{t-\tau}^i, X_t^j, \mathcal{S}$ )
15:      if  $|I| < I^{\min}(X_{t-\tau}^i \rightarrow X_t^j)$  then  $\triangleright$  Store minimum  $I$  of parent among all tests until now
16:         $I^{\min}(X_{t-\tau}^i \rightarrow X_t^j) = |I|$ 
17:        if p-value  $> \alpha$  then  $\triangleright$  Removed only after all  $X_{t-\tau}^i$  have been tested
18:          Mark  $X_{t-\tau}^i$  for removal from  $\tilde{\mathcal{P}}(X_t^j)$ 
19:          Break from inner while-loop
20:      Remove non-significant parents from  $\tilde{\mathcal{P}}(X_t^j)$ 
21:      Sort parents in  $\tilde{\mathcal{P}}(X_t^j)$  by  $I^{\min}(X_{t-\tau}^i \rightarrow X_t^j)$  from largest to smallest
22: return  $\tilde{\mathcal{P}}(X_t^j)$ 

```

Algorithm S2 Pseudo-code for MCI causal discovery step. In our analyses we set $p_Y = N\tau_{\max}$, that is, no restrictions on the maximum number of parents. Here we state the algorithm for $\tau \geq 0$, then causal links for $\tau = 0$ correspond to contemporaneous links, which are left undirected here.

Require: Time series dataset $\mathbf{X} = (X^1, X^2, \dots, X^N)$, significance level α , sorted parents $\mathcal{P}(X_t^j)$ for all variables X^j estimated with Algorithm S1, maximum time lag τ_{\max} , maximum number p_Y of parents of variable X^j , maximum number p_X of parents of variable X^i

```

1: for all  $(X_{t-\tau}^i \rightarrow X_t^j)$  with  $i = 1, \dots, N$  and  $\tau = 0, \dots, \tau_{\max}$ , excluding  $(X_t^j \rightarrow X_t^j)$  do
2:   Define  $\mathcal{P}_{p_Y}(X_t^j)$  as the first  $p_Y$  parents from  $\mathcal{P}(X_t^j)$  and remove  $X_{t-\tau}^i$  if necessary
3:   Define  $\mathcal{P}_{p_X}(X_{t-\tau}^i)$  as the first  $p_X$  parents from  $\mathcal{P}(X_{t-\tau}^i)$ , shifted by  $\tau$ 
4:   Run MCI test to obtain (p-value,  $I$ )  $\leftarrow$  CI( $X_{t-\tau}^i, X_t^j, \mathcal{S} = \{\mathcal{P}_{p_Y}(X_t^j), \mathcal{P}_{p_X}(X_{t-\tau}^i)\}$ )
5: Optionally adjust p-values of links  $(X_{t-\tau}^i \rightarrow X_t^j)$  by False Discovery Rate-approach with significance level  $\alpha$ 
6: return p-values or q-values (for FDR-adjusted tests) and MCI test statistic values

```

S9 Supplementary Tables

Table S3. Model coefficients for numerical experiments corresponding to different setups and link strengths $I_{X \rightarrow Y}$. (*Top table*) Coupling coefficients c for link $X \rightarrow Y$ in common driver model (Sect. S5) and all cross-links in the random models studied in Fig. 2. As described in Sect. S4, the link strength coefficient c for the dependencies $cf^{(1,2,3)}(x)$ is chosen such that in the corresponding single-link model with $X_t^i = \eta_t^i$ and $X_t^j = cf(X_{t-1}^i) + \eta_t^j$ we have a constant mutual information $I(\eta_{t-1}^j; cf(\eta_{t-1}^i) + \eta_t^j) = I_{X \rightarrow Y}$. Since no analytical results are available for all our coupling functions and noise types, this is achieved numerically by estimating mutual information with a state-of-the-art nearest-neighbor estimator [28] with nearest-neighbor parameter $k = 10$ (average of 500 realizations of sample length $T = 1000$). (*Bottom table*) Forcing coefficients b for different noise types, forcing strengths I , autocorrelations a and common drivers D_Z . As described in Sect. S5, the common driver forcing coefficient $b = b(D_Z, a)$ and the covariance among the drivers $\sigma_Z = b/2$ are chosen such that in the full model with no coupling $X \rightarrow Y$ for every pair ($D_Z > 0, a \geq 0$) with $c = 0$ we obtain the same unconditional dependency strength $I(X_{t-1}; Y_t)$.

Noise	Coupling function	Link strength $I_{X \rightarrow Y}$ [nats]				
		0.02	0.03	0.04	0.05	0.08
gauss	$f^{(1)}$	0.200	0.247	0.287	0.324	0.414
gauss	$f^{(2)}$	0.176	0.221	0.254	0.286	0.370
gauss	$f^{(3)}$	0.070	0.087	0.098	0.111	0.142
weibull	$f^{(1)}$	0.136	0.169	0.198	0.230	0.307
weibull	$f^{(2)}$	0.110	0.137	0.166	0.190	0.253
weibull	$f^{(3)}$	0.033	0.041	0.047	0.053	0.071

Noise	Forcing strength I [nats]	Autocorrelation a	Number of common drivers D_Z							
			1	2	3	4	8	12	16	32
gauss	0.1	0.00	0.894	0.558	0.431	0.358	0.233	0.182	0.153	0.100
gauss	0.1	0.40	0.692	0.445	0.346	0.291	0.191	0.150	0.126	0.082
gauss	0.1	0.60	0.488	0.320	0.254	0.211	0.142	0.112	0.094	0.062
gauss	0.1	0.80	0.242	0.164	0.131	0.110	0.077	0.060	0.053	0.035
gauss	0.1	0.90	0.112	0.078	0.063	0.054	0.038	0.031	0.026	0.018
gauss	0.1	0.95	0.046	0.032	0.026	0.023	0.017	0.013	0.012	0.008
gauss	0.4	0.00	1.802	1.035	0.782	0.642	0.409	0.316	0.263	0.169
gauss	0.4	0.40	1.402	0.835	0.637	0.524	0.337	0.261	0.217	0.140
gauss	0.4	0.60	0.986	0.614	0.471	0.395	0.255	0.198	0.166	0.108
gauss	0.4	0.80	0.502	0.331	0.260	0.218	0.145	0.115	0.097	0.064
gauss	0.4	0.90	0.252	0.169	0.136	0.117	0.079	0.063	0.053	0.036
gauss	0.4	0.95	0.120	0.082	0.068	0.058	0.040	0.032	0.028	0.019
gauss	0.8	0.00	3.226	1.686	1.246	1.018	0.638	0.487	0.403	0.256
gauss	0.8	0.40	2.526	1.372	1.025	0.842	0.530	0.405	0.337	0.215
gauss	0.8	0.60	1.775	1.024	0.772	0.635	0.404	0.313	0.260	0.167
gauss	0.8	0.80	0.902	0.570	0.438	0.367	0.239	0.187	0.157	0.102
gauss	0.8	0.90	0.459	0.302	0.238	0.201	0.135	0.106	0.089	0.059
gauss	0.8	0.95	0.230	0.155	0.123	0.106	0.073	0.058	0.050	0.034
weibull	0.4	0.00	1.043	0.640	0.492	0.412	0.266	0.206	0.172	0.112
weibull	0.4	0.40	0.832	0.522	0.404	0.338	0.221	0.172	0.145	0.094
weibull	0.4	0.60	0.591	0.382	0.299	0.250	0.167	0.130	0.110	0.072
weibull	0.4	0.80	0.297	0.201	0.162	0.136	0.093	0.073	0.063	0.042
weibull	0.4	0.90	0.144	0.100	0.079	0.069	0.046	0.039	0.032	0.022
weibull	0.4	0.95	0.063	0.042	0.036	0.030	0.021	0.017	0.015	0.011

Table S4. One-sided critical values for the ACE maximal correlation independence test statistic between two univariate samples. Note that the maximal correlation is non-negative. Since the variables are copula-transformed, these values are valid for any marginal distribution and depend only on the sample size T . In **Tigramite**, the whole null distribution for more sample sizes T is available as needed for FDR-pruning. Standard errors are estimated from a bootstrap with 1,000 surrogates.

Sample size	10%	5%	1%	0.1%	0.01%
50	0.5848 ± 0.0008	0.6287 ± 0.0014	0.7221 ± 0.0029	0.8325 ± 0.0018	0.8796 ± 0.0072
60	0.5395 ± 0.0009	0.5808 ± 0.0011	0.6714 ± 0.0025	0.8225 ± 0.0063	0.8453 ± 0.0113
70	0.4962 ± 0.0008	0.5364 ± 0.0012	0.6247 ± 0.0027	0.8180 ± 0.0112	0.8430 ± 0.0079
80	0.4628 ± 0.0008	0.5007 ± 0.0012	0.5834 ± 0.0018	0.7518 ± 0.0071	0.8388 ± 0.0072
90	0.4337 ± 0.0007	0.4696 ± 0.0010	0.5479 ± 0.0026	0.7372 ± 0.0182	0.8347 ± 0.0049
100	0.3940 ± 0.0007	0.4263 ± 0.0009	0.4927 ± 0.0018	0.5936 ± 0.0096	0.6938 ± 0.0173
150	0.3052 ± 0.0006	0.3297 ± 0.0007	0.3773 ± 0.0014	0.4378 ± 0.0032	0.5153 ± 0.0503
200	0.2562 ± 0.0004	0.2768 ± 0.0006	0.3194 ± 0.0010	0.3687 ± 0.0029	0.4144 ± 0.0157
250	0.2215 ± 0.0004	0.2401 ± 0.0005	0.2763 ± 0.0009	0.3187 ± 0.0028	0.3630 ± 0.0093
300	0.1980 ± 0.0004	0.2148 ± 0.0005	0.2469 ± 0.0007	0.2887 ± 0.0022	0.3194 ± 0.0069
350	0.1804 ± 0.0003	0.1953 ± 0.0004	0.2247 ± 0.0006	0.2558 ± 0.0023	0.2948 ± 0.0114
400	0.1662 ± 0.0003	0.1806 ± 0.0004	0.2075 ± 0.0006	0.2392 ± 0.0018	0.2714 ± 0.0027
450	0.1550 ± 0.0003	0.1684 ± 0.0004	0.1941 ± 0.0006	0.2249 ± 0.0021	0.2514 ± 0.0052
500	0.1456 ± 0.0003	0.1578 ± 0.0003	0.1820 ± 0.0006	0.2100 ± 0.0016	0.2353 ± 0.0046
550	0.1383 ± 0.0002	0.1501 ± 0.0003	0.1736 ± 0.0006	0.2006 ± 0.0009	0.2259 ± 0.0028
600	0.1313 ± 0.0003	0.1431 ± 0.0003	0.1651 ± 0.0006	0.1908 ± 0.0017	0.2205 ± 0.0053
650	0.1253 ± 0.0002	0.1361 ± 0.0003	0.1568 ± 0.0006	0.1809 ± 0.0010	0.1994 ± 0.0031
700	0.1199 ± 0.0002	0.1304 ± 0.0003	0.1502 ± 0.0006	0.1747 ± 0.0011	0.1976 ± 0.0059
750	0.1157 ± 0.0002	0.1257 ± 0.0003	0.1444 ± 0.0005	0.1647 ± 0.0013	0.1883 ± 0.0047
800	0.1118 ± 0.0002	0.1216 ± 0.0002	0.1406 ± 0.0004	0.1622 ± 0.0013	0.1842 ± 0.0044
850	0.1081 ± 0.0002	0.1174 ± 0.0003	0.1356 ± 0.0004	0.1563 ± 0.0007	0.1724 ± 0.0028
900	0.1044 ± 0.0002	0.1134 ± 0.0003	0.1304 ± 0.0004	0.1528 ± 0.0011	0.1674 ± 0.0013
950	0.1014 ± 0.0002	0.1104 ± 0.0002	0.1274 ± 0.0005	0.1469 ± 0.0008	0.1650 ± 0.0037
1,000	0.0986 ± 0.0002	0.1072 ± 0.0002	0.1240 ± 0.0005	0.1423 ± 0.0015	0.1612 ± 0.0035
1,250	0.0873 ± 0.0001	0.0949 ± 0.0002	0.1102 ± 0.0004	0.1274 ± 0.0012	0.1419 ± 0.0051
1,500	0.0796 ± 0.0002	0.0863 ± 0.0002	0.0996 ± 0.0003	0.1143 ± 0.0009	0.1319 ± 0.0024
1,750	0.0730 ± 0.0001	0.0795 ± 0.0002	0.0921 ± 0.0004	0.1079 ± 0.0009	0.1222 ± 0.0035
2,000	0.0681 ± 0.0001	0.0741 ± 0.0002	0.0860 ± 0.0003	0.0994 ± 0.0005	0.1111 ± 0.0022
2,500	0.0605 ± 0.0001	0.0660 ± 0.0001	0.0765 ± 0.0003	0.0890 ± 0.0008	0.1004 ± 0.0032
3,000	0.0550 ± 0.0001	0.0600 ± 0.0001	0.0692 ± 0.0002	0.0822 ± 0.0007	0.0918 ± 0.0013
3,500	0.0509 ± 0.0001	0.0553 ± 0.0001	0.0643 ± 0.0002	0.0743 ± 0.0006	0.0844 ± 0.0011
4,000	0.0474 ± 0.0001	0.0515 ± 0.0001	0.0598 ± 0.0002	0.0694 ± 0.0005	0.0767 ± 0.0007
5,000	0.0422 ± 0.0001	0.0460 ± 0.0001	0.0534 ± 0.0002	0.0622 ± 0.0005	0.0683 ± 0.0008
6,000	0.0385 ± 0.0001	0.0420 ± 0.0001	0.0490 ± 0.0001	0.0571 ± 0.0005	0.0650 ± 0.0012
7,000	0.0354 ± 0.0001	0.0386 ± 0.0001	0.0446 ± 0.0001	0.0511 ± 0.0004	0.0567 ± 0.0009
8,000	0.0331 ± 0.0001	0.0361 ± 0.0001	0.0422 ± 0.0001	0.0491 ± 0.0004	0.0560 ± 0.0008
9,000	0.0313 ± 0.0001	0.0341 ± 0.0001	0.0396 ± 0.0001	0.0461 ± 0.0004	0.0500 ± 0.0008
10,000	0.0296 ± 0.0001	0.0324 ± 0.0001	0.0378 ± 0.0001	0.0436 ± 0.0002	0.0480 ± 0.0008
12,000	0.0271 ± 0.0001	0.0296 ± 0.0001	0.0344 ± 0.0001	0.0397 ± 0.0003	0.0449 ± 0.0010
14,000	0.0249 ± 0.0000	0.0272 ± 0.0001	0.0316 ± 0.0001	0.0365 ± 0.0003	0.0413 ± 0.0010
16,000	0.0233 ± 0.0000	0.0255 ± 0.0001	0.0296 ± 0.0001	0.0343 ± 0.0003	0.0393 ± 0.0010
18,000	0.0220 ± 0.0000	0.0240 ± 0.0000	0.0278 ± 0.0001	0.0327 ± 0.0003	0.0372 ± 0.0005
20,000	0.0208 ± 0.0000	0.0227 ± 0.0000	0.0265 ± 0.0001	0.0310 ± 0.0003	0.0348 ± 0.0002
25,000	0.0186 ± 0.0000	0.0203 ± 0.0000	0.0237 ± 0.0001	0.0275 ± 0.0003	0.0319 ± 0.0010
30,000	0.0170 ± 0.0000	0.0186 ± 0.0000	0.0216 ± 0.0001	0.0248 ± 0.0002	0.0289 ± 0.0006
40,000	0.0147 ± 0.0000	0.0160 ± 0.0000	0.0187 ± 0.0001	0.0217 ± 0.0002	0.0239 ± 0.0003
50,000	0.0131 ± 0.0000	0.0143 ± 0.0000	0.0167 ± 0.0001	0.0195 ± 0.0002	0.0217 ± 0.0003

Table S5. Spurious and causal drivers of the North Atlantic Oscillation component No. 5 in the Earth system application. A link $X_{t-\tau}^i \rightarrow X_t^j$ is denoted as component No. $i \rightarrow_\tau$ No. j . The time resolution and therefore the time delay τ is given in units of three days. The q-value denotes the FDR-corrected p-value taking into account all cross- and autodependency links in a network.

Causal links	Correlation (p-/q-value)	bivMCI (p-/q-value)	GC (p-/q-value)	MCI ₀ (p-/q-value)	MCI (p-/q-value)
2→ ₁ 5	-0.05 (0.102 / 0.277)	-0.20 (0.000 / 0.000)	-0.10 (0.018 / 0.429)	-0.12 (0.000 / 0.012)	-0.16 (0.000 / 0.000)
55→ ₂ 5	0.13 (0.000 / 0.001)	0.14 (0.000 / 0.001)	0.13 (0.003 / 0.202)	0.15 (0.000 / 0.001)	0.14 (0.000 / 0.002)
44→ ₁ 5	-0.02 (0.531 / 0.730)	-0.09 (0.006 / 0.116)	-0.13 (0.002 / 0.162)	-0.13 (0.000 / 0.009)	-0.14 (0.000 / 0.004)
43→ ₁ 5	0.08 (0.012 / 0.066)	0.12 (0.000 / 0.014)	0.08 (0.058 / 0.628)	0.09 (0.004 / 0.135)	0.12 (0.000 / 0.020)
45→ ₁ 5	-0.28 (0.000 / 0.000)	-0.13 (0.000 / 0.005)	-0.09 (0.031 / 0.519)	-0.10 (0.002 / 0.091)	-0.11 (0.001 / 0.036)
Spurious links					
8→ ₇ 5	-0.15 (0.000 / 0.000)	-0.06 (0.075 / 0.436)	-0.04 (0.330 / 0.886)	-0.02 (0.521 / 0.949)	0.00 (0.953 / 0.996)
9→ ₅ 5	0.13 (0.000 / 0.001)	0.07 (0.026 / 0.265)	0.07 (0.097 / 0.712)	0.00 (0.965 / 0.999)	0.04 (0.245 / 0.845)
9→ ₃ 5	0.12 (0.000 / 0.003)	0.08 (0.009 / 0.147)	0.05 (0.239 / 0.839)	0.02 (0.506 / 0.945)	0.05 (0.117 / 0.698)
9→ ₄ 5	0.12 (0.000 / 0.003)	-0.00 (0.933 / 0.987)	-0.02 (0.709 / 0.972)	-0.03 (0.370 / 0.905)	-0.04 (0.273 / 0.858)
8→ ₆ 5	-0.12 (0.000 / 0.003)	-0.06 (0.063 / 0.404)	-0.01 (0.761 / 0.980)	-0.07 (0.025 / 0.372)	-0.05 (0.134 / 0.725)
22→ ₆ 5	-0.12 (0.000 / 0.004)	-0.07 (0.041 / 0.335)	-0.05 (0.201 / 0.822)	-0.04 (0.219 / 0.824)	-0.04 (0.195 / 0.802)
9→ ₇ 5	0.12 (0.000 / 0.004)	0.03 (0.369 / 0.784)	0.10 (0.014 / 0.393)	0.06 (0.084 / 0.630)	0.02 (0.568 / 0.952)
31→ ₄ 5	-0.11 (0.000 / 0.005)	-0.05 (0.091 / 0.472)	-0.05 (0.250 / 0.845)	-0.06 (0.071 / 0.594)	-0.07 (0.041 / 0.479)
22→ ₅ 5	-0.11 (0.001 / 0.008)	-0.04 (0.229 / 0.675)	-0.03 (0.551 / 0.949)	0.01 (0.777 / 0.987)	0.01 (0.819 / 0.986)
38→ ₁ 5	0.11 (0.001 / 0.008)	0.03 (0.403 / 0.804)	-0.00 (0.939 / 0.993)	0.03 (0.395 / 0.913)	0.01 (0.662 / 0.972)
15→ ₇ 5	0.11 (0.001 / 0.009)	0.02 (0.444 / 0.825)	0.04 (0.303 / 0.877)	0.07 (0.044 / 0.491)	0.04 (0.193 / 0.800)

Table S6. Strongest causal links according to MCI statistic in the Earth system application. A link $X_{t-\tau}^i \rightarrow X_t^j$ is denoted as component No. $i \rightarrow_\tau j$. The time resolution and therefore the time delay τ is given in units of three days. The q-value denotes the FDR-corrected p-value taking into account all cross- and autodependency links in a network for $\tau > 0$.

Causal links	Correlation (p-/q-value)	bivMCI (p-/q-value)	GC (p-/q-value)	MCI ₀ (p-/q-value)	MCI (p-/q-value)
30→ ₁ 15	-0.17 (0.000 / 0.000)	-0.22 (0.000 / 0.000)	-0.23 (0.000 / 0.000)	-0.25 (0.000 / 0.000)	-0.25 (0.000 / 0.000)
55→ ₁ 41	-0.09 (0.007 / 0.044)	-0.21 (0.000 / 0.000)	-0.15 (0.001 / 0.087)	-0.23 (0.000 / 0.000)	-0.24 (0.000 / 0.000)
15→ ₁ 2	0.24 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.15 (0.000 / 0.061)	0.23 (0.000 / 0.000)	0.22 (0.000 / 0.000)
20→ ₁ 16	0.16 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.18 (0.000 / 0.010)	0.22 (0.000 / 0.000)	0.21 (0.000 / 0.000)
7→ ₁ 0	0.30 (0.000 / 0.000)	0.21 (0.000 / 0.000)	0.12 (0.006 / 0.279)	0.24 (0.000 / 0.000)	0.21 (0.000 / 0.000)
35→ ₁ 4	-0.21 (0.000 / 0.000)	-0.19 (0.000 / 0.000)	-0.17 (0.000 / 0.014)	-0.21 (0.000 / 0.000)	-0.21 (0.000 / 0.000)
47→ ₁ 49	-0.22 (0.000 / 0.000)	-0.24 (0.000 / 0.000)	-0.15 (0.001 / 0.090)	-0.18 (0.000 / 0.000)	-0.20 (0.000 / 0.000)
22→ ₁ 10	-0.15 (0.000 / 0.000)	-0.22 (0.000 / 0.000)	-0.17 (0.000 / 0.012)	-0.19 (0.000 / 0.000)	-0.20 (0.000 / 0.000)
42→ ₁ 29	0.08 (0.013 / 0.070)	-0.13 (0.000 / 0.005)	-0.19 (0.000 / 0.003)	-0.22 (0.000 / 0.000)	-0.20 (0.000 / 0.000)
10→ ₁ 38	-0.12 (0.000 / 0.002)	-0.15 (0.000 / 0.000)	-0.18 (0.000 / 0.008)	-0.20 (0.000 / 0.000)	-0.20 (0.000 / 0.000)
43→ ₁ 2	0.23 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.18 (0.000 / 0.008)	0.21 (0.000 / 0.000)	0.20 (0.000 / 0.000)
55→ ₁ 1	0.10 (0.003 / 0.022)	0.11 (0.001 / 0.021)	0.12 (0.006 / 0.287)	0.20 (0.000 / 0.000)	0.20 (0.000 / 0.000)
51→ ₁ 18	-0.24 (0.000 / 0.000)	-0.18 (0.000 / 0.000)	-0.15 (0.000 / 0.069)	-0.19 (0.000 / 0.000)	-0.19 (0.000 / 0.000)
20→ ₁ 48	-0.14 (0.000 / 0.000)	-0.18 (0.000 / 0.000)	-0.20 (0.000 / 0.001)	-0.18 (0.000 / 0.000)	-0.19 (0.000 / 0.000)
40→ ₁ 41	-0.17 (0.000 / 0.000)	-0.22 (0.000 / 0.000)	-0.15 (0.001 / 0.087)	-0.20 (0.000 / 0.000)	-0.19 (0.000 / 0.000)
27→ ₁ 32	0.20 (0.000 / 0.000)	0.21 (0.000 / 0.000)	0.21 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.19 (0.000 / 0.000)
18→ ₁ 59	-0.29 (0.000 / 0.000)	-0.23 (0.000 / 0.000)	-0.09 (0.027 / 0.490)	-0.19 (0.000 / 0.000)	-0.19 (0.000 / 0.000)
48→ ₁ 35	0.24 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.13 (0.002 / 0.178)	0.14 (0.000 / 0.002)	0.19 (0.000 / 0.000)
33→ ₁ 27	-0.23 (0.000 / 0.000)	-0.15 (0.000 / 0.000)	-0.13 (0.002 / 0.159)	-0.19 (0.000 / 0.000)	-0.19 (0.000 / 0.000)
35→ ₁ 1	-0.15 (0.000 / 0.000)	-0.21 (0.000 / 0.000)	-0.19 (0.000 / 0.002)	-0.18 (0.000 / 0.000)	-0.18 (0.000 / 0.000)
40→ ₁ 7	0.29 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.12 (0.005 / 0.276)	0.18 (0.000 / 0.000)	0.18 (0.000 / 0.000)
20→ ₁ 38	-0.19 (0.000 / 0.000)	-0.20 (0.000 / 0.000)	-0.13 (0.002 / 0.175)	-0.13 (0.000 / 0.011)	-0.18 (0.000 / 0.000)
35→ ₁ 48	-0.06 (0.086 / 0.249)	-0.15 (0.000 / 0.001)	-0.18 (0.000 / 0.005)	-0.18 (0.000 / 0.000)	-0.18 (0.000 / 0.000)
23→ ₁ 0	-0.05 (0.094 / 0.264)	-0.18 (0.000 / 0.000)	-0.18 (0.000 / 0.010)	-0.13 (0.000 / 0.005)	-0.18 (0.000 / 0.000)
21→ ₁ 46	-0.24 (0.000 / 0.000)	-0.24 (0.000 / 0.000)	-0.17 (0.000 / 0.021)	-0.18 (0.000 / 0.000)	-0.18 (0.000 / 0.000)
5→ ₁ 2	-0.07 (0.033 / 0.134)	-0.13 (0.000 / 0.005)	-0.18 (0.000 / 0.008)	-0.12 (0.000 / 0.028)	-0.17 (0.000 / 0.000)
9→ ₁ 6	-0.18 (0.000 / 0.000)	-0.20 (0.000 / 0.000)	-0.18 (0.000 / 0.004)	-0.16 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
4→ ₂ 52	-0.20 (0.000 / 0.000)	-0.17 (0.000 / 0.000)	-0.04 (0.380 / 0.907)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
6→ ₁ 29	-0.17 (0.000 / 0.000)	-0.15 (0.000 / 0.000)	-0.13 (0.003 / 0.219)	-0.16 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
8→ ₁ 6	0.08 (0.019 / 0.091)	0.17 (0.000 / 0.000)	0.12 (0.006 / 0.296)	0.12 (0.000 / 0.017)	0.17 (0.000 / 0.000)
7→ ₁ 23	0.15 (0.000 / 0.000)	0.20 (0.000 / 0.000)	0.08 (0.074 / 0.669)	0.15 (0.000 / 0.001)	0.17 (0.000 / 0.000)
1→ ₁ 56	-0.23 (0.000 / 0.000)	-0.21 (0.000 / 0.000)	-0.17 (0.000 / 0.022)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
48→ ₂ 6	-0.16 (0.000 / 0.000)	-0.18 (0.000 / 0.000)	-0.04 (0.313 / 0.880)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
6→ ₁ 49	0.22 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.13 (0.002 / 0.155)	0.18 (0.000 / 0.000)	0.17 (0.000 / 0.000)
27→ ₁ 22	0.13 (0.000 / 0.001)	0.21 (0.000 / 0.000)	0.17 (0.000 / 0.024)	0.17 (0.000 / 0.000)	0.17 (0.000 / 0.000)
9→ ₁ 24	0.20 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.10 (0.014 / 0.396)	0.17 (0.000 / 0.000)	0.17 (0.000 / 0.000)
14→ ₁ 43	-0.17 (0.000 / 0.000)	-0.15 (0.000 / 0.000)	-0.20 (0.000 / 0.001)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
52→ ₁ 0	-0.14 (0.000 / 0.000)	-0.14 (0.000 / 0.002)	-0.19 (0.000 / 0.002)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
1→ ₁ 48	0.28 (0.000 / 0.000)	0.14 (0.000 / 0.002)	0.15 (0.001 / 0.089)	0.22 (0.000 / 0.000)	0.17 (0.000 / 0.000)
24→ ₁ 43	-0.10 (0.002 / 0.021)	-0.14 (0.000 / 0.001)	-0.15 (0.001 / 0.089)	-0.18 (0.000 / 0.000)	-0.17 (0.000 / 0.000)
43→ ₁ 23	-0.15 (0.000 / 0.000)	-0.21 (0.000 / 0.000)	-0.16 (0.000 / 0.027)	-0.15 (0.000 / 0.001)	-0.17 (0.000 / 0.000)
51→ ₁ 9	0.30 (0.000 / 0.000)	0.17 (0.000 / 0.000)	0.16 (0.000 / 0.028)	0.18 (0.000 / 0.000)	0.16 (0.000 / 0.000)
8→ ₁ 35	0.17 (0.000 / 0.000)	0.16 (0.000 / 0.000)	0.11 (0.013 / 0.387)	0.18 (0.000 / 0.000)	0.16 (0.000 / 0.000)
40→ ₁ 1	0.15 (0.000 / 0.000)	0.12 (0.000 / 0.011)	0.17 (0.000 / 0.017)	0.12 (0.000 / 0.013)	0.16 (0.000 / 0.000)
25→ ₁ 39	0.16 (0.000 / 0.000)	0.17 (0.000 / 0.000)	0.13 (0.003 / 0.216)	0.17 (0.000 / 0.000)	0.16 (0.000 / 0.000)
56→ ₁ 26	0.19 (0.000 / 0.000)	0.17 (0.000 / 0.000)	0.13 (0.002 / 0.170)	0.17 (0.000 / 0.000)	0.16 (0.000 / 0.000)
29→ ₁ 52	-0.15 (0.000 / 0.000)	-0.18 (0.000 / 0.000)	-0.15 (0.001 / 0.089)	-0.15 (0.000 / 0.001)	-0.16 (0.000 / 0.000)
7→ ₁ 26	0.20 (0.000 / 0.000)	0.15 (0.000 / 0.000)	0.14 (0.001 / 0.107)	0.19 (0.000 / 0.000)	0.16 (0.000 / 0.000)
9→ ₁ 18	-0.29 (0.000 / 0.000)	-0.18 (0.000 / 0.000)	-0.14 (0.001 / 0.150)	-0.12 (0.000 / 0.014)	-0.16 (0.000 / 0.000)
28→ ₁ 6	-0.12 (0.000 / 0.002)	-0.12 (0.000 / 0.008)	-0.11 (0.007 / 0.301)	-0.18 (0.000 / 0.000)	-0.16 (0.000 / 0.000)
11→ ₁ 24	-0.22 (0.000 / 0.000)	-0.19 (0.000 / 0.000)	-0.13 (0.002 / 0.177)	-0.17 (0.000 / 0.000)	-0.16 (0.000 / 0.000)

Table S7. As in Tab. S6, but for contemporaneous links at lag $\tau = 0$, which are left undirected here. Note that FDR adjustment was conducted only among all contemporaneous links. Since we do not assess causality among contemporaneous interactions, and only excluded lagged effects, many of these links are likely spurious correlations.

Causal links	Correlation (p-/q-value)	bivMCI (p-/q-value)	GC (p-/q-value)	MCI ₀ (p-/q-value)	MCI (p-/q-value)
2-26	0.63 (0.000 / 0.000)	0.60 (0.000 / 0.000)	0.52 (0.000 / 0.000)	0.51 (0.000 / 0.000)	0.53 (0.000 / 0.000)
26-2	0.63 (0.000 / 0.000)	0.60 (0.000 / 0.000)	0.52 (0.000 / 0.000)	0.44 (0.000 / 0.000)	0.51 (0.000 / 0.000)
18-2	-0.53 (0.000 / 0.000)	-0.50 (0.000 / 0.000)	-0.44 (0.000 / 0.000)	-0.45 (0.000 / 0.000)	-0.45 (0.000 / 0.000)
2-18	-0.53 (0.000 / 0.000)	-0.50 (0.000 / 0.000)	-0.44 (0.000 / 0.000)	-0.44 (0.000 / 0.000)	-0.44 (0.000 / 0.000)
26-42	0.56 (0.000 / 0.000)	0.51 (0.000 / 0.000)	0.42 (0.000 / 0.000)	0.39 (0.000 / 0.000)	0.43 (0.000 / 0.000)
42-26	0.56 (0.000 / 0.000)	0.51 (0.000 / 0.000)	0.42 (0.000 / 0.000)	0.43 (0.000 / 0.000)	0.42 (0.000 / 0.000)
32-18	0.50 (0.000 / 0.000)	0.46 (0.000 / 0.000)	0.43 (0.000 / 0.000)	0.34 (0.000 / 0.000)	0.39 (0.000 / 0.000)
18-32	0.50 (0.000 / 0.000)	0.46 (0.000 / 0.000)	0.43 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.39 (0.000 / 0.000)
4-2	0.31 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.33 (0.000 / 0.000)
4-7	0.27 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.33 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.31 (0.000 / 0.000)
1-4	0.20 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.31 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.30 (0.000 / 0.000)
1-6	0.30 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.32 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.30 (0.000 / 0.000)
7-4	0.27 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.33 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.30 (0.000 / 0.000)
33-0	0.50 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.29 (0.000 / 0.000)
6-1	0.30 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.32 (0.000 / 0.000)	0.20 (0.000 / 0.000)	0.29 (0.000 / 0.000)
26-4	0.41 (0.000 / 0.000)	0.38 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.29 (0.000 / 0.000)
2-4	0.31 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.28 (0.000 / 0.000)
4-26	0.41 (0.000 / 0.000)	0.38 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.28 (0.000 / 0.000)	0.28 (0.000 / 0.000)
16-0	0.36 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.28 (0.000 / 0.000)	0.27 (0.000 / 0.000)	0.27 (0.000 / 0.000)
2-42	0.39 (0.000 / 0.000)	0.33 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.27 (0.000 / 0.000)
0-33	0.50 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.27 (0.000 / 0.000)
0-16	0.36 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.28 (0.000 / 0.000)	0.15 (0.000 / 0.000)	0.26 (0.000 / 0.000)
1-11	0.24 (0.000 / 0.000)	0.28 (0.000 / 0.000)	0.34 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.26 (0.000 / 0.000)
1-18	0.41 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.25 (0.000 / 0.000)
32-2	-0.32 (0.000 / 0.000)	-0.28 (0.000 / 0.000)	-0.29 (0.000 / 0.000)	-0.24 (0.000 / 0.000)	-0.25 (0.000 / 0.000)
22-2	-0.07 (0.042 / 0.095)	-0.24 (0.000 / 0.000)	-0.25 (0.000 / 0.000)	-0.23 (0.000 / 0.000)	-0.25 (0.000 / 0.000)
4-1	0.20 (0.000 / 0.000)	0.35 (0.000 / 0.000)	0.31 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.25 (0.000 / 0.000)
42-2	0.39 (0.000 / 0.000)	0.33 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.25 (0.000 / 0.000)
0-13	0.32 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.24 (0.000 / 0.000)
33-22	0.34 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.24 (0.000 / 0.000)
26-18	-0.28 (0.000 / 0.000)	-0.29 (0.000 / 0.000)	-0.22 (0.000 / 0.000)	-0.21 (0.000 / 0.000)	-0.24 (0.000 / 0.000)
18-16	0.19 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.16 (0.000 / 0.000)	0.24 (0.000 / 0.000)
16-18	0.19 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.24 (0.000 / 0.000)
1-7	0.30 (0.000 / 0.000)	0.31 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.24 (0.000 / 0.000)
22-33	0.34 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.23 (0.000 / 0.000)	0.24 (0.000 / 0.000)
13-0	0.32 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.24 (0.000 / 0.000)
18-26	-0.28 (0.000 / 0.000)	-0.29 (0.000 / 0.000)	-0.22 (0.000 / 0.000)	-0.27 (0.000 / 0.000)	-0.24 (0.000 / 0.000)
22-18	0.29 (0.000 / 0.000)	0.27 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.24 (0.000 / 0.000)
40-7	0.26 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.24 (0.000 / 0.000)
18-22	0.29 (0.000 / 0.000)	0.27 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.23 (0.000 / 0.000)
7-40	0.26 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.21 (0.000 / 0.000)	0.23 (0.000 / 0.000)
16-22	0.30 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.23 (0.000 / 0.000)	0.23 (0.000 / 0.000)
18-1	0.41 (0.000 / 0.000)	0.36 (0.000 / 0.000)	0.26 (0.000 / 0.000)	0.23 (0.000 / 0.000)	0.23 (0.000 / 0.000)
11-1	0.24 (0.000 / 0.000)	0.28 (0.000 / 0.000)	0.34 (0.000 / 0.000)	0.20 (0.000 / 0.000)	0.23 (0.000 / 0.000)
7-1	0.30 (0.000 / 0.000)	0.31 (0.000 / 0.000)	0.30 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.23 (0.000 / 0.000)
2-22	-0.07 (0.042 / 0.095)	-0.24 (0.000 / 0.000)	-0.25 (0.000 / 0.000)	-0.23 (0.000 / 0.000)	-0.23 (0.000 / 0.000)
2-32	-0.32 (0.000 / 0.000)	-0.28 (0.000 / 0.000)	-0.29 (0.000 / 0.000)	-0.20 (0.000 / 0.000)	-0.22 (0.000 / 0.000)
18-28	0.23 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.22 (0.000 / 0.000)
27-1	0.20 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.24 (0.000 / 0.000)	0.18 (0.000 / 0.000)	0.22 (0.000 / 0.000)
22-16	0.30 (0.000 / 0.000)	0.29 (0.000 / 0.000)	0.25 (0.000 / 0.000)	0.19 (0.000 / 0.000)	0.22 (0.000 / 0.000)
34-33	0.14 (0.000 / 0.000)	0.22 (0.000 / 0.000)	0.17 (0.000 / 0.002)	0.18 (0.000 / 0.000)	0.22 (0.000 / 0.000)

S10 Supplementary Figures

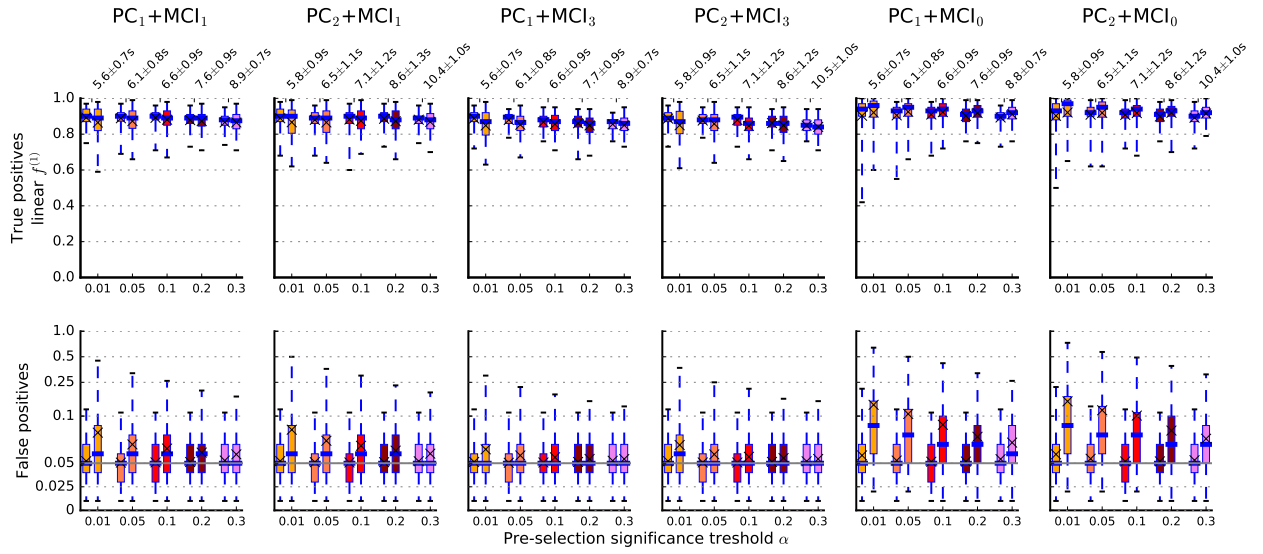


Figure S4. Influence of different PC thresholds α and number of tested condition combinations q_{\max} in the condition-selection Algorithm S1 as well as number of conditions p_X in Algorithm S2 on the performance of PCMCi. We do not limit the maximum condition dimension p_{\max} . Here the network model with linear links for ParCorr with $N = 20$ and $T = 150$ was used. We ran ten different coupling topologies with 100 realizations each of weakly and strongly autocorrelated variables at a link strength of $I_{X \rightarrow Y} = 0.04$. The pairs of boxplots (1%, 25%, 75%, 99% percentiles, median marked by a bar and mean marked by 'x') always denote weakly (left) and strongly (right) autocorrelated variable pairs. The bottom row shows the distribution of FPRs and the upper row the distributions of TPRs for different thresholds α increasing along the x -axis in each plot. Note the logarithmic y -axis in the bottom panel for $\text{FPR} > 0.1$. Average runtime and standard deviation are given on top. See discussion in Sect. S1.




Figure S5. Same as Fig. S4, but for GPACE. Here the network model with nonlinear links with $N = 10$ and $T = 250$ was used.




Figure S6. More detailed results for the numerical experiments from Fig. 2A,B demonstrating the detection biases in the ParCorr implementation. The network model as described in Sect. S4 is here run with $N = 20$ and $T = 150$. The pairs of boxplots (1%, 25%, 75%, 99% percentiles, median marked by a bar and mean marked by ‘x’) always denote weakly (left) and strongly (right) autocorrelated variable pairs. The bottom row shows the distribution of FPRs and the upper row the distributions of TPRs for different link strengths increasing along the x -axis in each plot. Note the logarithmic y -axis in the bottom panel for $FPR > 0.1$. The more confined a distribution, the less biased is the TPR. GC and PC are compared with PCMCi for different numbers of additional MCI conditions: PC₁+MCI for $p_X = 1$ and PC₁+MCI₃ for $p_X = 3$. Also PC₁+MCI₀ with $p_X = 0$ is shown. PC(raw) refers to the PC algorithm without pruning of links. See Sect. S4 for a discussion.




Figure S7. More detailed results for the numerical experiments from Fig. 2C demonstrating detection biases in the GPACE implementation for the network model described in Sect. S4 with $N = 5$ and $T = 250$. For GPACE we here show results of Link-TPRs for all types of nonlinear links $f^{(2)}(x) = (1 - 4e^{-x^2/2})x$ and $f^{(3)}(x) = (1 - 4x^3e^{-x^2/2})x$. GC and PC are compared with PCMCI for different numbers of additional MCI conditions: PC_1+MCI for $p_X = 1$ and PC_1+MCI_3 for $p_X = 3$. Also PC_1+MCI_0 with $p_X = 0$ is shown. PC(raw) refers to the PC algorithm without pruning of links. See Sect. S4 for a discussion.




Figure S8. Same as in Fig. S7 but with $T = 500$ samples. As expected, then PCMCi has higher power also for nonlinear dependencies.




Figure S9. More detailed results for the numerical experiments from Fig. 3A in the ParCorr implementation where the performance of different causal discovery methods for larger network sizes is studied. Here the link strength was fixed to $I_{X \rightarrow Y} = 0.04$ and 10 different coupling topologies were tested to limit computational time. The sample length is $T = 150$. The pairs of boxplots (1%, 25%, 75%, 99% percentiles, median marked by a bar and mean marked by 'x') always denote weakly (left) and strongly (right) autocorrelated variable pairs. The bottom row shows the distribution of FPRs and the upper row the distributions of TPRs for different network sizes N increasing along the x -axis in each plot. Average runtime and standard deviation are given on top. See Sect. S4 for a discussion.

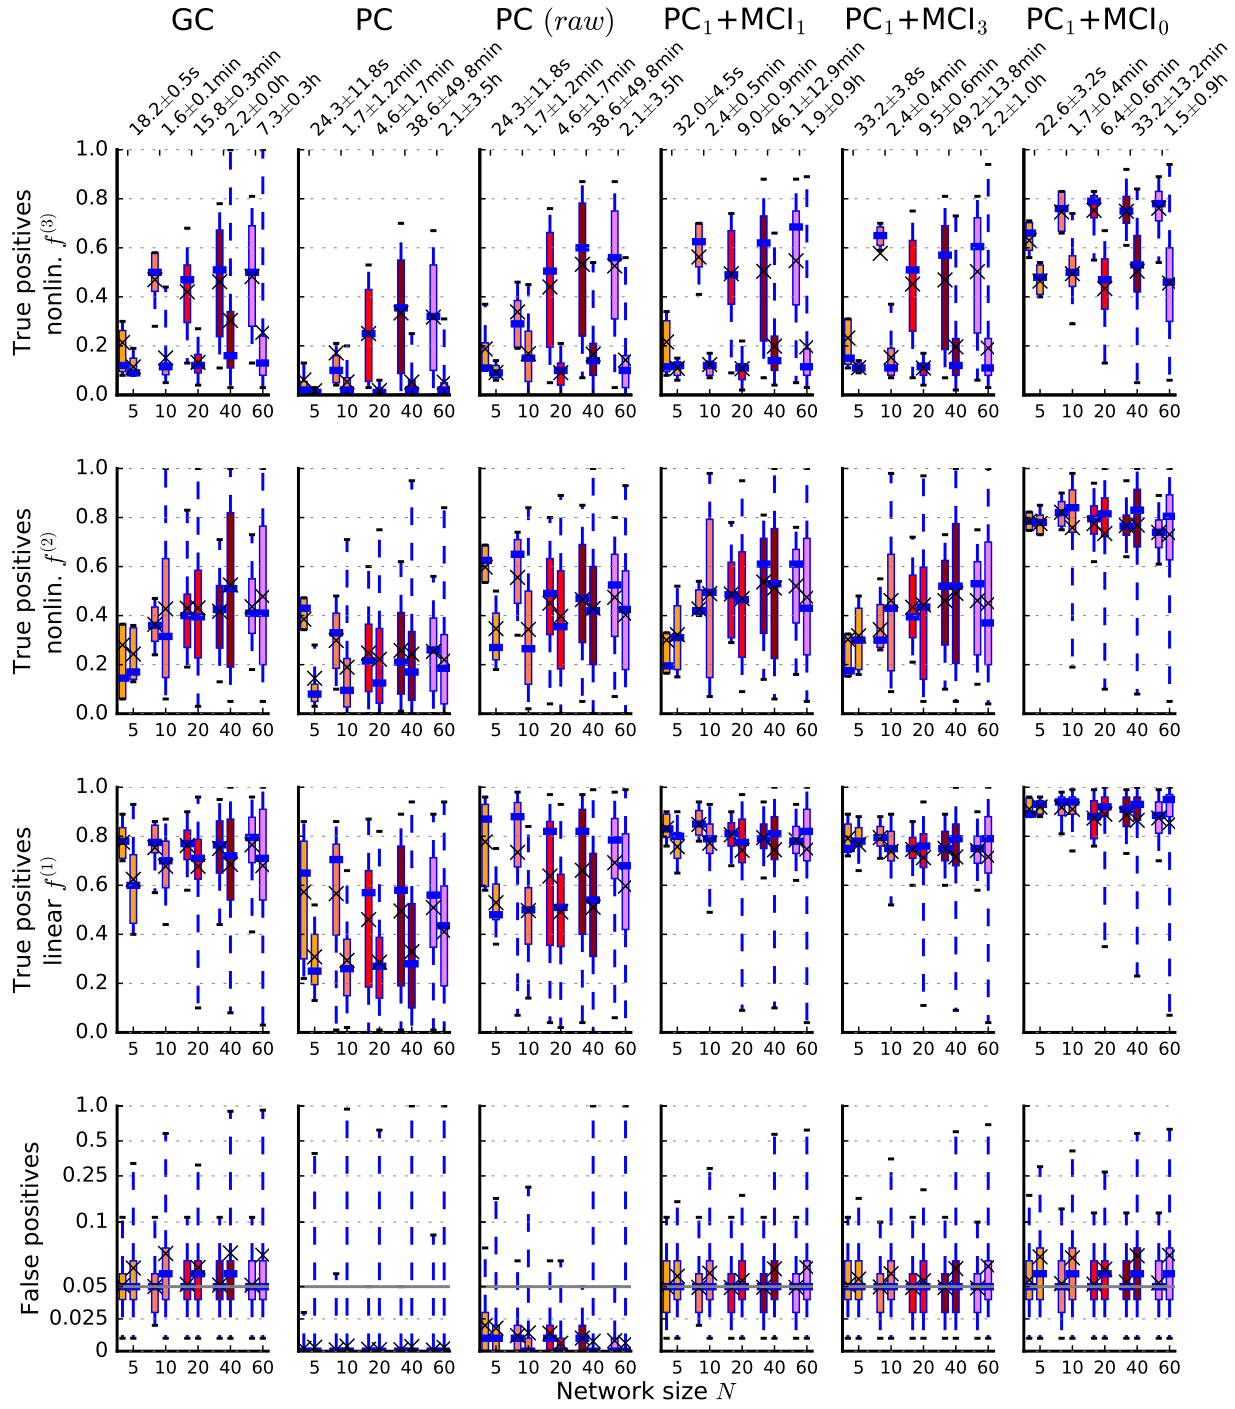


Figure S10. Same as in Fig. S9 but for GPACE implementation (Fig. 3B) with $T = 250$. We here show results of TPRs for all types of nonlinear links $f^{(2)}(x) = (1 - 4e^{-x^2/2})x$ and $f^{(3)}(x) = (1 - 4x^3e^{-x^2/2})x$.

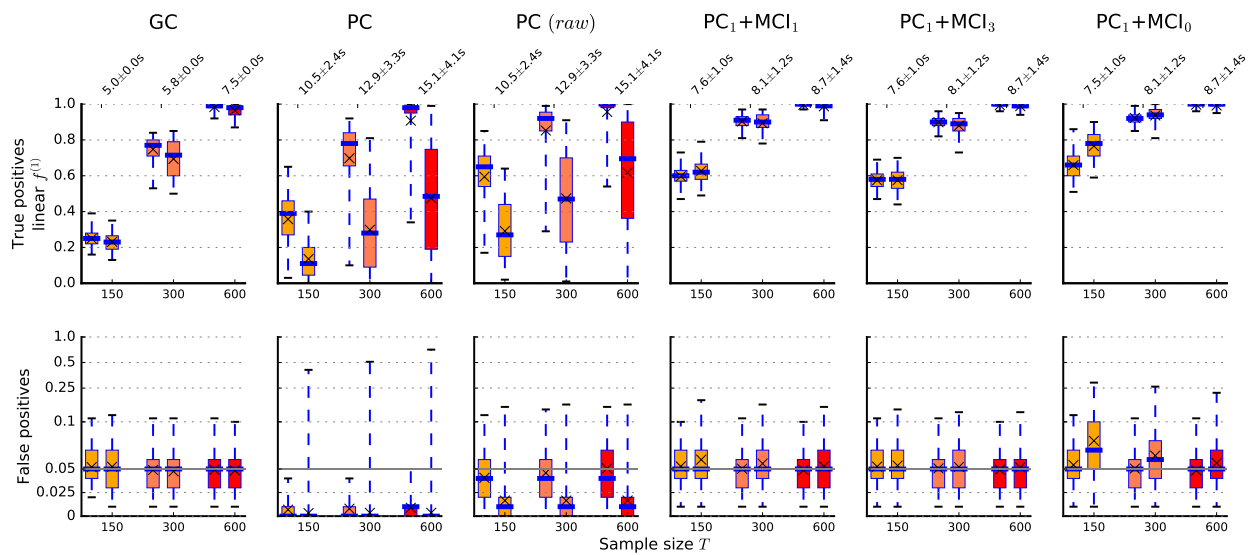


Figure S11. More detailed results for the numerical experiments for larger sample sizes for the ParCorr implementation of the network model described in Sect. S4 with $N = 20$. Here the link strength was fixed to $I_{X \rightarrow Y} = 0.02$ and 10 different coupling topologies were tested to limit computational time. PC(raw) refers to the PC algorithm without pruning of links.

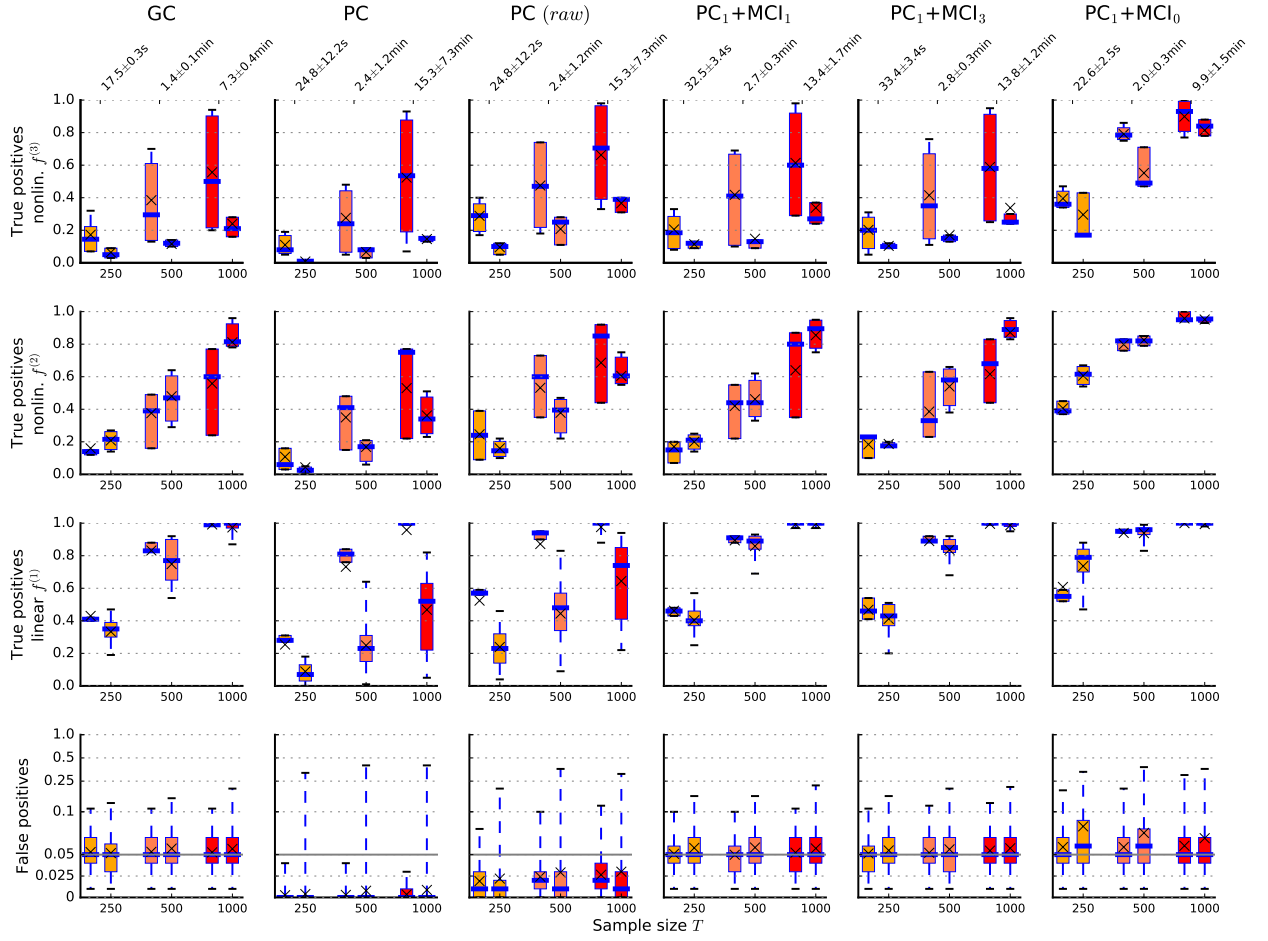


Figure S12. Same as in Fig. S11 for GPACE implementation for $N = 5$. For GPACE we here show results of Link-TPRs for all types of nonlinear links $f^{(2)}(x) = (1 - 4e^{-x^2/2})x$ and $f^{(3)}(x) = (1 - 4x^3e^{-x^2/2})x$. PC-MCI has higher power levels than GC throughout. Note the steep increase in computation time since GPACE scales $\sim T^3$.




Figure S13. As in Fig. S9 for ParCorr, but for a variant of the model where the noise terms η in the model described in Sect. S4 are spatially correlated with a covariance of $\sigma_{ij} = 0.2$ for $i, j \in \{1, \dots, N\}$. Here the link strength was fixed to $I_{X \rightarrow Y} = 0.02$ and 10 different coupling topologies were tested to limit computational time. The sample size is doubled here to $T = 300$. Here we show results for the condition-selection algorithm run with threshold $\alpha = 0.2$ and in Fig. S14 for $\alpha = 0.4$. PC₁+MCI_{all} refers to the MCI test where *all* estimated parents of X are used, that is, $p_X = |\mathcal{P}(X_{t-\tau})|$.

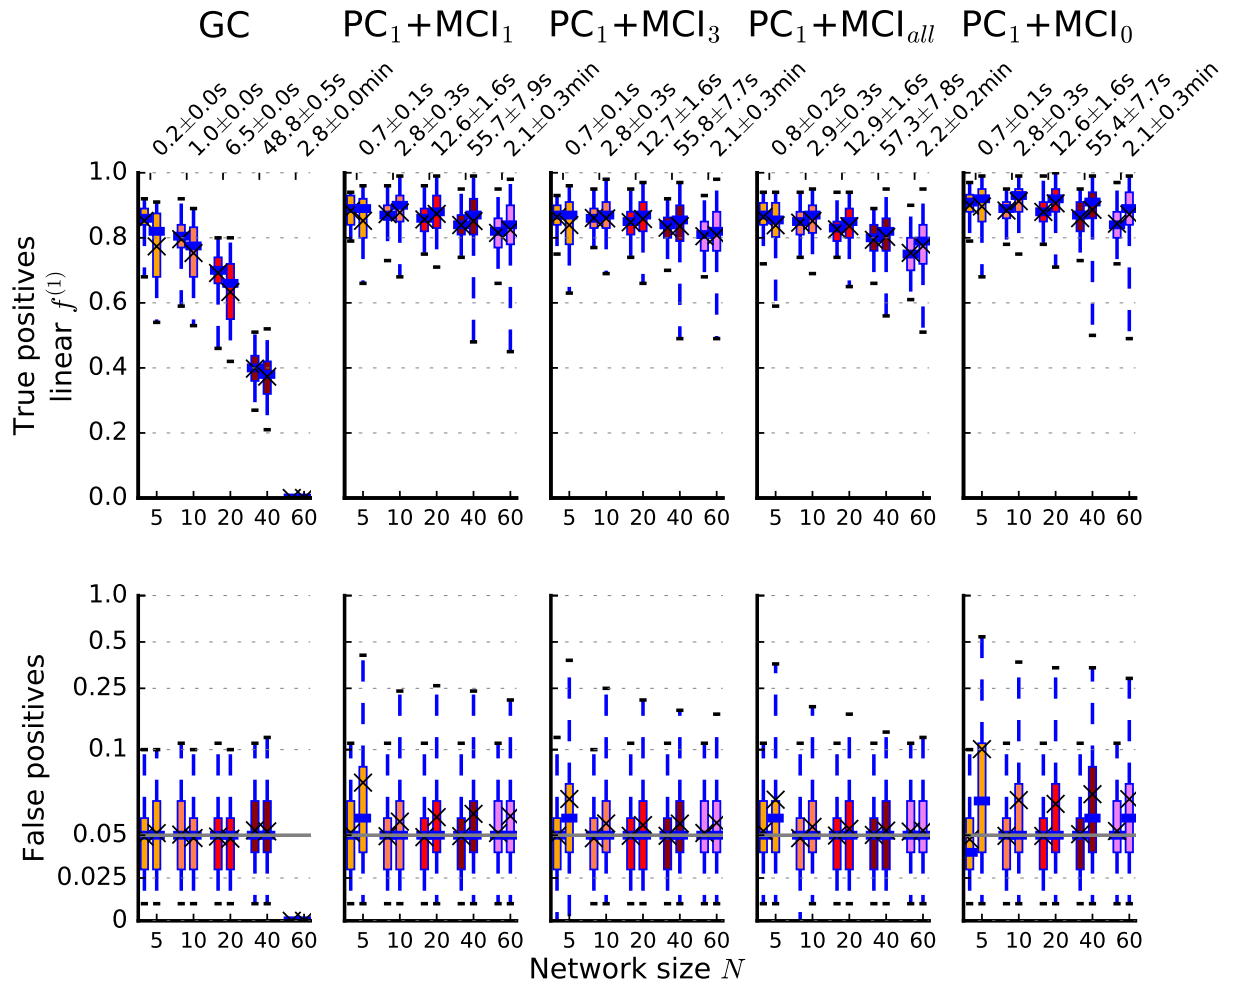


Figure S14. As in Fig. S13, but for $\alpha = 0.4$ which better restricts FPRs.

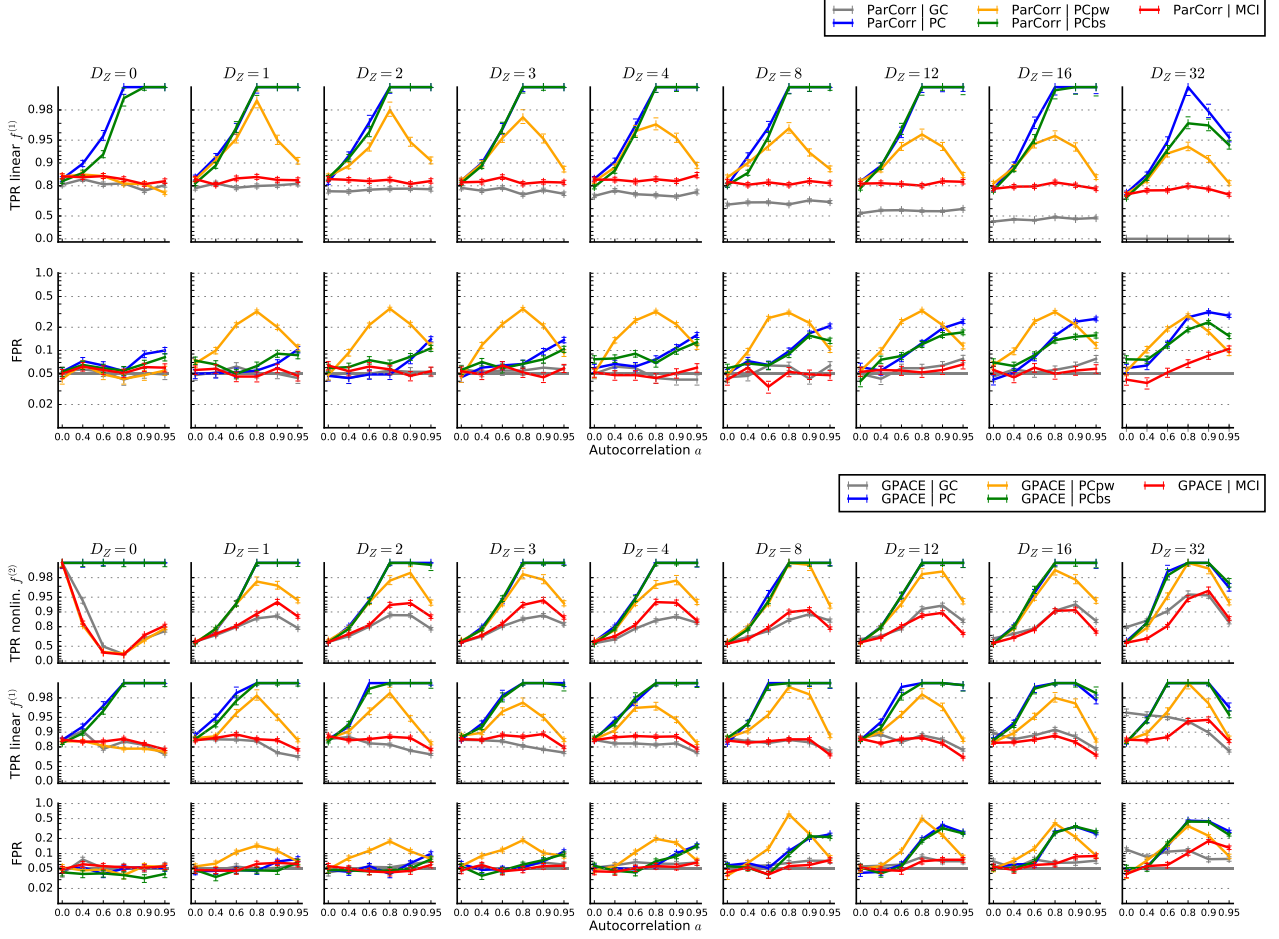


Figure S15. Numerical experiments with simple common driver model for ParCorr (top panel) and GPACE (bottom panel). The model time series graph is depicted in Fig. S3, the parameters are described in Sect. S5 and Tab. S3. The link strength is the same for all autocorrelations and dimensions. Next to the comparison of the GC, PC, and MCI tests, we also show alternative autocorrelation remedies via pre-whitening (PCpw) and a computationally expensive block-shuffle surrogate test (PCbs) described in Sect. S5. For $D_Z = 0$ the model setup corresponds to two autocorrelated processes without a common driver forcing. The sample lengths are $T = 150$ for ParCorr and $T = 250$ for GPACE. 1,000 realizations were evaluated to assess false and true positives. The test significance level is 95% implying an expected FPR of 5% (gray line). Note the logarithmic y-axis (inverted logarithm for TPR) and the custom x-axis. The error bars provide bootstrapped standard errors. For $D_Z = 32$ the dimensionality of the GC test is 171 and exceeds the time series length for ParCorr. While the ParCorr GC test then cannot be evaluated anymore with traditional estimators, GPACE shows still some performance in such high dimensions. See discussion in Sect. S5.




Figure S16. Further GPAC numerical experiments with simple common driver model for Weibull noise terms discussed in Sect. S5. Here we choose a coupling strength of $I_{X \rightarrow Y} = 0.05$ for the linear dependency and $I_{X \rightarrow Y} = 0.08$ for the nonlinear one. See discussion in Sect. S5.

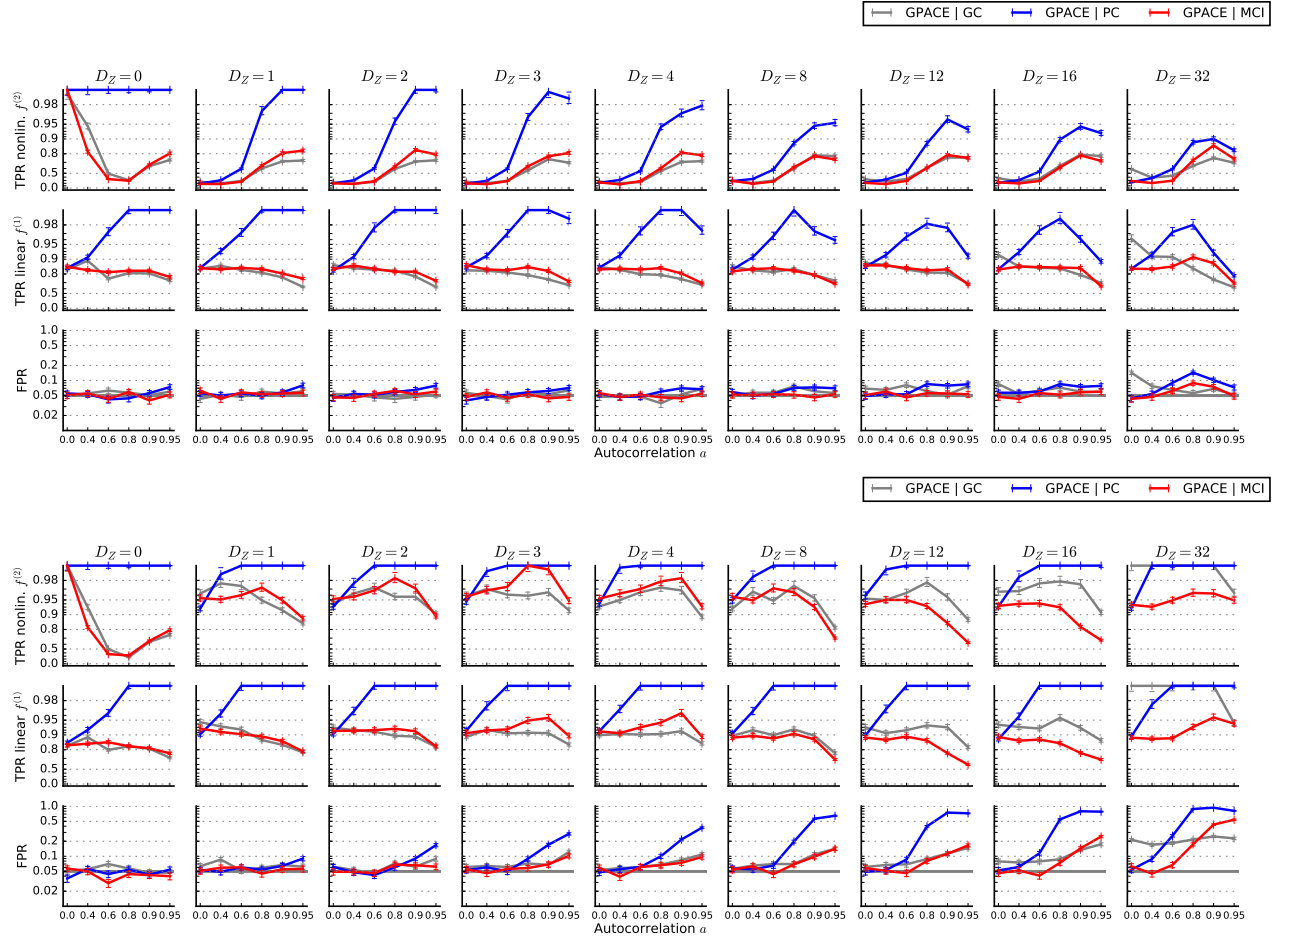


Figure S17. Further GPACE numerical experiments with simple common driver model for different forcing strengths I of the common drivers Z , that is, different coefficients b (see Tab. S3), corresponding to an unconditional correlation of 0.4 ($I = 0.1$, top) and 0.9 ($I = 0.8$ bottom) discussed in Sect. S5. See discussion in Sect. S5.




Figure S18. Same as Fig. S15, but for $T = 300$ samples for ParCorr (top) and $T = 500$ samples for GPACE (bottom). We choose a coupling strength of $I_{X \rightarrow Y} = 0.02$ for the linear dependency and $I_{X \rightarrow Y} = 0.05$ for the nonlinear one for GPACE and $I_{X \rightarrow Y} = 0.02$ for ParCorr. See discussion in Sect. S5. Note that for $D_Z > 2$ the inverted logarithmic scale makes the only slight increase in power of about 5% look larger.




Figure S19. Comparison of different measures for Earth system example for lagged links with $\tau > 0$. The value of MCI at the lag with maximum absolute MCI value is compared with Pearson correlation, bivariate MCI, Granger causality, and MCI_0 (also at their corresponding lags with maximum absolute value). For details on the analysis see Sect. S7. While Correlation or bivariate MCI and MCI tests correspond to very different underlying independence hypotheses, the MCI vs. MCI_0 scatter plot is naturally more similar since the only difference lies in the additional condition on the strongest parent of the supposed driving variable. Nevertheless, for individual links, this can lead to strong differences in line with our numerical experiments. The marked links are the strongest in each measure and some links discussed in the North Atlantic Oscillation example. In Tab. S6, we provide p-values and FDR-adjusted q-values for the strongest links.

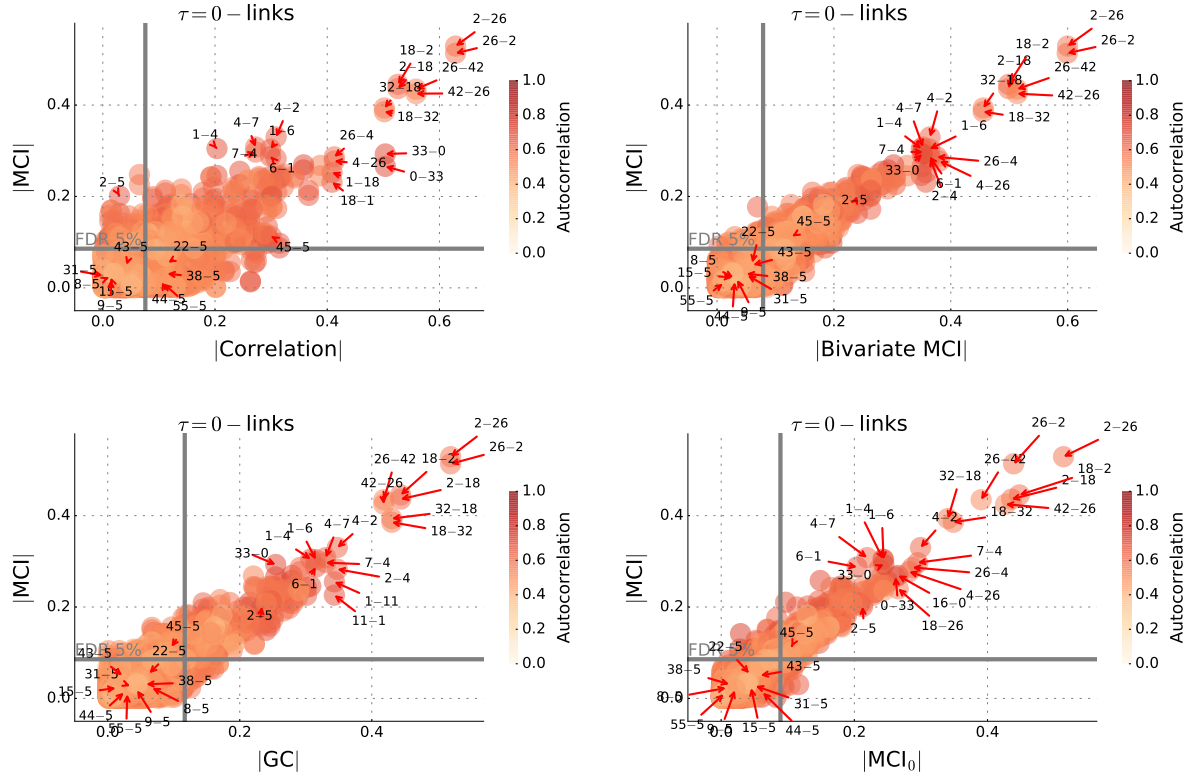


Figure S20. As in Fig. S19, but for contemporaneous links at lag $\tau = 0$, which are left undirected here. Note that FDR adjustment was conducted only among all contemporaneous links. In Tab. S7 we provide p-values and FDR-adjusted q-values for the strongest links. Since we do not assess causality among contemporaneous interactions, and only excluded lagged effects, many of these links are likely spurious correlations.




Figure S21. Robustness of MCI link strength for different PC_1 parameters α .




Figure S22. Component regions for strongest causal links as labeled in Fig. S19. In Tab. S6, we provide p-values and FDR-adjusted q-values. Here only the core region is depicted, in [58] complete region loadings are shown.




Figure S23. Analysis of North Atlantic Oscillation component No. 5. (**Top left**) Drivers (parents) of No. 5. Contours lines of different colors identify the core region loadings belonging to each causal driver component (note the dipole of the NAO, for all regions see [47]). The link color gives the MCI value, and the label denotes the time lag in units of 3 days (if multiple time lags are involved, the strongest is listed first). P-values and FDR-adjusted q-values are given in Tab. S5. (**Bottom left**) Components driven by No. 5. (**Top right**) Contemporaneous links of No. 5 according to significance of MCI test $X_t^i \perp\!\!\!\perp X_t^5 \mid \mathcal{P}(X_t^5), \mathcal{P}_{p_X=3}(X_t^i)$. Directionality here can either be resolved by a finer, but noisier, time-resolution or with recently developed methods that address the inference of causality if no time-ordering is given at all [32, 38, 52]. Since we do not assess causality among contemporaneous interactions, and only excluded lagged effects, many of these links are likely spurious correlations. (**Bottom right**) Contemporaneous links of No. 5 according to significance of MCI test $X_t^5 \perp\!\!\!\perp X_t^j \mid \mathcal{P}(X_t^j), \mathcal{P}_{p_X=3}(X_t^5)$. Since only the p_X strongest parents are conditioned on, the two tests are slightly different.

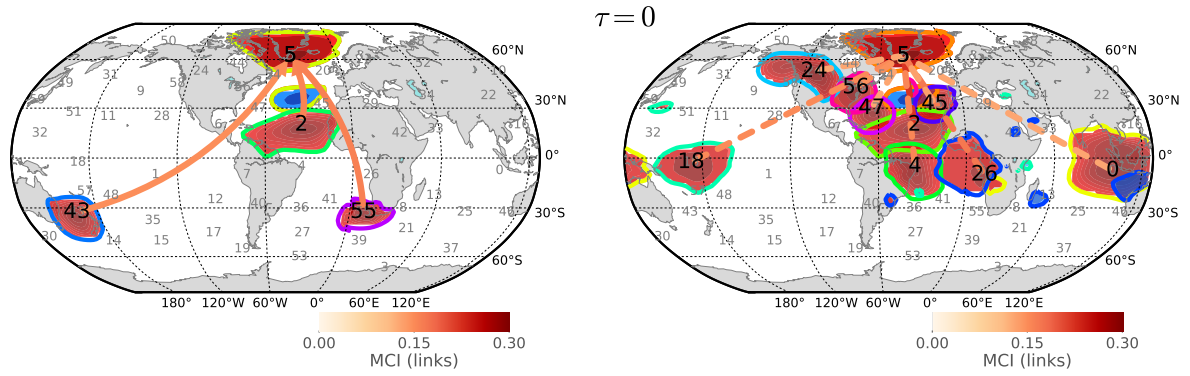


Figure S24. Same as Fig. S23 top plots, but for GPACE.